

DECOMPOSITION OF THE 605 KW BASE-TO-PEAK LOAD GAP IN A HOSPITAL BUILDING: EVIDENCE OF OCCUPANCY SCHEDULE DOMINANCE OVER COOLING LOAD IN DEMAND-SIDE MANAGEMENT STRATEGIES

Radimas Putra Muhammad Davi Labib*

Electronics Engineering Education Program,
Universitas Negeri Jakarta,
Indonesia

Bagus Tri Kuncoro

Electronics Engineering Education Program,
Universitas Negeri Jakarta,
Indonesia

Ade Ayu Rahmawati

Electronics Engineering Education Program,
Universitas Negeri Jakarta,
Indonesia

Vina Oktaviani

Electronics Engineering Education Program,
Universitas Negeri Jakarta,
Indonesia

Agam Nizar Dwi Nur Fahmi

Electronics Engineering Education Program,
Universitas Negeri Jakarta,
Indonesia

Made Dika Nugraha

Computer Engineering Program,
Universitas Warmadewa,
Indonesia

Article Info

Article History:

Received: 6 November 2025

Revised: 22 November 2025

Accepted: 10 December 2025

DOI : 10.21009/jvote.v8i2.72074



Keywords:

Decomposition
Load factor
Percentile clustering
Load profile
Hospital

Abstract

The reliability of a hospital's power supply hinges on the accurate determination of backup source capacity; however, standard design practices typically rely on static installed power values without accounting for dynamic load fluctuations. This study characterizes the electrical load profile of a healthcare facility using a reference load profile derived from building energy simulations, comprising 8,760 hourly observations over a full year. The analysis procedure includes timestamp validation, power balance verification, end-use decomposition, descriptive statistics for weekday and weekend profiles, percentile-based clustering to distinguish between base load (5th percentile) and peak load (95th percentile) regimes, and vulnerability zone mapping via load duration curves. The recorded average load was 1,156.3 kW, with an annual peak of 1,576.2 kW and a load factor of 0.734. The dynamic range between the base load (880.2 kW) and peak load (1,484.9 kW) was 604.7 kW (38.4% of the peak load), while the non-sheddable vital load measured 695.7 kW. Decomposition of this range revealed a decoupling between energy share and peak contribution: cooling systems accounted for 43.3% of annual energy consumption but contributed only 10.4% to the dynamic range, whereas lighting and interior equipment collectively contributed 59.7%. The weekend load deficit was only 14.7% (177.4 kW) and was concentrated during active hours; the difference was less than 1% during nighttime hours but reached 26.9% at 16:00. Vulnerability mapping indicates that a 17.5% reduction in generator capacity is achievable through a maximum load shed of 276.2 kW, all of which falls within the category of deferrable loads. These findings shift demand-side management priorities toward scheduled loads rather than air conditioning systems.

To cite this article : Labib, R. P. M. D., Kuncoro, B. T., Rahmawati, A. A., Oktaviani, V., Fahmi, A. N. D. N., & Nugraha, M. D. (2025). Decomposition of the 605 kW Base-to-Peak Load Gap in a Hospital Building: Evidence of Occupancy Schedule Dominance over Cooling Load in Demand-Side Management Strategies. *Jurnal Pendidikan Vokasional Teknik Elektronika*, 8(2), 91-101

INTRODUCTION

Among all building typologies, hospitals place the highest demands on the reliability of their power supply. This is due not merely to the volume of consumption, but to the nature of their operations, which run continuously—24 hours a day, seven days a week—without the holiday cycles

Corresponding author:

Radimas P. M. D. L., Universitas Negeri Jakarta, Indonesia (RadimasPutraMuhammadDaviLabib@unj.ac.id)

typical of offices or commercial buildings (Soyler & Izgi, 2022). The loads served also vary in criticality: operating rooms, intensive care units, neonatal incubators, ventilators, diagnostic imaging equipment, and sterile-room HVAC systems cannot tolerate supply interruptions; consequently, power outages in healthcare facilities pose immediate risks to patient safety and the continuity of medical services (Farquharson et al., 2018; Ibiam et al., 2025). To meet these demands, hospital electrical systems are designed with multiple layers: generator sets (gensets) serve as backup sources, Uninterruptible Power Supply (UPS) units act as transitional buffers, and systems have recently evolved into microgrid architectures that integrate generators, photovoltaics, and battery energy storage (Ali et al., 2023; Saki et al., 2025). However, the actual reliability achieved by these layers is determined not simply by their presence, but by how precisely their capacities are matched to the characteristics of the loads they must serve.

This is where the fundamental issue arises. Backup system design practices typically rely on a single static figure—such as installed capacity or total connected load multiplied by a normative demand factor—without considering how the actual load fluctuates over time. The consequences are well-documented. A study of 23 public hospitals in Istanbul with over 100 beds found that actual peak demand was only about one-third of the installed system capacity; consequently, the installed transformers and generators were nearly three times larger than actual needs, resulting in wasteful initial investment as well as ongoing operation and maintenance costs throughout the system's lifespan (Soyler & Izgi, 2022). Similar trends have long been noted in hospital distribution system design, where capacity is often determined based on connected load rather than measured demand (Neal, 2013), whereas approaches grounded in load profiles and outage scenarios have been shown to yield far more cost-effective generator and battery configurations (Dong et al., 2018; Leon et al., 2023). The problem is twofold: setting capacity without understanding the magnitude and timing of load spikes also creates a risk of overloading the generator during power transfer. Both types of failure stem from the same limitation: a single figure can never adequately represent a load distribution that fluctuates on an hourly basis.

While literature on electrical loads in healthcare buildings has expanded, the primary focus has been on predictive modeling rather than structural characterization. Healthcare facilities have been studied less frequently than commercial, office, hotel, and residential buildings—largely due to the complexity of their consumption patterns; despite the increasing availability of load data driven by advancements in metering technology, detailed analyses based on measured data remain limited (Cao et al., 2020). Most research has concentrated on short-term load forecasting—utilizing single or ensemble machine learning models (Cao et al., 2020) and artificial neural networks for large-scale facilities (Bagnasco et al., 2014; Bagnasco et al., 2015; Bertolini et al., 2013; Morinigo-Sotelo et al., 2011)—as well as on the analysis and forecasting of intra-day consumption (Gordillo-Orquera et al., 2018). Other studies have examined hospitals within the frameworks of energy management and consumption intensity, covering topics ranging from aggregate costs and consumption (Hu et al., 2004) and facility energy audits and management (Alhurayess & Darwish, 2012) to specific end-uses like vertical transportation systems (Katsanis et al., 2011) and national-scale building stock assessments (Balaras et al., 2007). The outputs of such studies typically consist of prediction accuracy metrics or intensity figures rather than a structural breakdown of the load itself. Consequently, a significant research gap remains: few studies have granularly analyzed 8,760-hour time-series data to disentangle persistent base load components from transient peak load components and subsequently attribute them to specific end-uses. Yet, this specific attribution is a prerequisite for designing load-shedding schemes, as decisions regarding load shedding require knowledge of which loads fluctuate, the magnitude of those fluctuations, and the times at which they occur.

Addressing this gap, this study aims to conduct a descriptive analysis and percentile-based clustering of hourly active power time-series data—collected over a full year from a hospital building—decomposed by end-use categories: HVAC loads (cooling and fans), interior lighting, and interior equipment (comprising both medical and support devices). Percentile clustering is employed to statistically define the base-load regime (5th percentile) and peak-load regime (95th percentile), allowing the difference between the two to be broken down into the quantitative contribution of each end-use. The results of this decomposition serve as the basis for establishing load priority levels for the design of a SCADA-based smart load-shedding scheme, adhering to the

smart-load demand-side management framework (Palensky & Dietrich, 2011) and the principle of tiered medical load classification based on criticality (Ibiam et al., 2025). It should be noted at the outset that this study is limited to formulating the analytical basis; the implementation and testing of control devices fall outside its scope.

The study's novelty lies in a counter-intuitive finding: the decoupling of energy share from the contribution to peak demand. Cooling loads account for 43.3% of annual energy consumption yet contribute only 10.4% to the 604.7 kW dynamic range spanning the 880.2 kW base load and the 1,484.9 kW peak load; conversely, lighting and interior equipment collectively account for 59.7% of that range. This finding corrects a long-held design assumption—that HVAC systems are the primary drivers of load variability—thereby shifting priorities for load-shedding interventions and reserve capacity allocation toward targets different from those typically assumed. The article is structured as follows: The Methods section outlines data sources, preprocessing and temporal validation procedures, and load characterization metrics. The Results and Discussion section presents annual statistical characteristics, daily and weekly profiles, monthly dynamics, the base-to-peak range decomposition, and implications for load-shedding design. The Conclusion section summarizes the quantitative findings and prospects for future development.

METHODS

Research Design

This study employs a quantitative, descriptive-analytical approach with an **ex post facto** design, utilizing secondary time-series data. The object of analysis is the active power profile of a healthcare facility building, recorded at hourly resolution over a full calendar year. This design was selected because the research objective is structural characterization—specifically, the separation and attribution of load components—rather than causal hypothesis testing or predictive modeling. The procedure comprises four sequential stages: data acquisition, preprocessing, statistical analysis and percentile clustering, and the mapping of outage vulnerability zones.

Data Sources and Structure

The data analyzed is secondary data consisting of an hourly-resolution reference load profile for a hospital building, comprising 8,760 observations representing 365 days of full operation. The dataset includes 11 variables: a timestamp (Date/Time), one total facility power channel (Electricity: Facility), five electricity channels categorized by end-use (Cooling, Fans, InteriorLights, InteriorEquipment, and Heating: Electricity), and four gas-based thermal energy channels. The analysis is limited to the electricity channels, aligning with the study's focus on power distribution system loading and backup source capacity. The Heating: Electricity channel was excluded after verification revealed zero values across all 8,760 data points, indicating that the heating function of the subject building is powered entirely by gas energy.

One point must be stated unequivocally: the profile used is the output of a building energy simulation based on a reference hospital typology—following the commercial sector reference building model framework (Deru et al., 2011)—rather than the result of field measurements using power meters. The methodological implications are significant: the profile is deterministic regarding the modeled occupancy schedule, contains no instrumentation noise, and does not reflect operational anomalies such as patient surges or equipment failures. These limitations are revisited in the Results and Discussion section.

Data Preprocessing

Reconstruction and validation of dating

The database timestamps contain only the month-day-hour format, omitting the year component. Consequently, the day-of-the-week attribute is undefined and must be reconstructed based on calendar assumptions. Since the comparison between weekdays and weekends is highly sensitive to these assumptions, all seven possible calendar configurations were systematically tested. The selection criterion was the internal consistency of the weekly rhythm—specifically, the emergence of two consecutive low-load days recurring on a seven-day cycle. The assumption satisfying this criterion was then applied to the entire temporal analysis. The hourly index was

normalized to a 0–23 range using the "hour-ending" convention, such that the 01:00:00 timestamp represents the 00:00–01:00 interval.

Power balance consistency checking

Data completeness was checked for missing values, duplicate timestamps, and negative values. Subsequently, submeter coverage was assessed by examining the power balance, defining the residual component as the difference between the facility's total power and the sum of the available end-use channels:

$$P_{res}(t) = P_{total}(t) - \sum_i P_i(t) \quad (1)$$

This residual component represents loads that are not explicitly sub-metered—such as pumps, heat rejection systems, exterior lighting, and hot water systems. Rather than being omitted, this component is treated as a distinct category to ensure the power balance remains closed and no magnitude is lost during the decomposition.

Determination of the load category taxonomy

Load categorization is not approached in a binary manner—distinguishing simply between critical and non-critical—at the channel level. The reason is straightforward: every end-use channel in a hospital comprises both persistent and schedule-dependent components, so labeling at the channel level would obscure the distinction between the two. Instead, each channel is first broken down into its base (floor) and variable components:

$$P_{i,floor} = \min_t P_i(t), t = 1, \dots, 8760 \quad (2)$$

$$P_{i,var}(t) = P_i(t) - P_{i,floor} \quad (3)$$

The base component represents a constant power demand throughout the year, whereas the variable component represents demand dependent on occupancy and operating conditions. Based on this distinction, loads are classified into three priority levels, as presented in Table 1.

Table 1. Taxonomy of load priority levels based on persistence and function

Level	Category	Components	Determination basis	Treatment
P1	Vital load	Floor components of all channels, including residuals	Persistence over 8,760 hours of observation	Non-sheddable
P2	Essential load	Cooling channel and fan variable components	Ventilation, differential pressure, and air change rate functions for sterile rooms	Modulatable, unbreakable
P3	Deferred load	Variable components of the InteriorLights, InteriorEquipment, and residual channels	Adherence to occupancy schedules and non-emergency activities	Sheddable/shiftable

This taxonomy serves as a proxy and is explicitly stated as a research assumption, as the available sub-meter granularity does not allow for the separation of vital medical equipment loads from administrative and support equipment loads sharing the same "InteriorEquipment" channel.

Analysis Techniques

Descriptive statistics and daily profile

Initial characterization was performed by calculating measures of central tendency and dispersion across all channels, including the mean, minimum, maximum, standard deviation, coefficient of variation, and load factor.

$$CV = \frac{\sigma}{\mu} \times 100\% \quad (4)$$

$$LF = \frac{P_{avg}}{P_{max}} \quad (5)$$

Typical daily load profiles are constructed by aggregating conditional averages based on hourly indices—categorized into weekdays (Monday–Friday) and weekends (Saturday–Sunday)—resulting in two 24-point profile curves. The statistical significance of the mean difference between these two

categories is tested, and the relative magnitude is expressed as the weekend deficit ratio. Monthly dynamics are examined through monthly quartile distributions visualized via boxplots, as well as by tracking the monthly trajectories of base, average, and peak loads.

Percentile clustering

Loading regimes are separated using percentile-based clustering—rather than quartile-based—to isolate extreme regimes relevant to capacity determination within a sufficiently narrow range. The p -th percentile of the total power distribution is denoted as Q_p , such that the set of hours for each regime is defined as:

$$H_{base} = \{ t \mid P_{total}(t) \leq Q_{0,05} \} \quad (6)$$

$$H_{peak} = \{ t \mid P_{total}(t) \geq Q_{0,95} \} \quad (7)$$

The dynamic range between regimes—the core metric in this analysis—is calculated as the difference between the two percentiles:

$$\Delta P = Q_{0,95} - Q_{0,05} \quad (8)$$

The contribution of each end-use to that dynamic range is obtained through the decomposition of the difference in conditional means between the two sets of regimes:

$$\kappa_i = \frac{P_{i,peak} - P_{i,base}}{P_{total,peak} - P_{total,base}} \times 100\% \quad (9)$$

The notations $P_{i,peak}$ and $P_{i,base}$ represent the conditional averages for channel i within the sets H_{peak} and H_{base} , respectively. The value of κ_i addresses the attribution question that lies at the core of the study's novelty: specifically, which end-use truly drives load fluctuations, regardless of its annual energy share.

Mapping of blackout vulnerability zones

Vulnerability zones are operationalized using a load duration curve, constructed by arranging 8,760 total power values in descending order. For each candidate backup source capacity C , three vulnerability indicators are defined: peak shedding magnitude, energy to be shed, and exposure duration:

$$S(C) = \max_t [P_{total}(t) - C]^+ \quad (10)$$

$$E_{shed}(C) = \sum_t [P_{total}(t) - C]^+ \cdot \Delta t \quad (11)$$

$$\rho(C) = \frac{n\{t \mid P_{total}(t) > C\}}{8760} \times 100\% \quad (12)$$

The three indicators are evaluated across several capacity scenarios to determine whether the deficit in each scenario can be covered by the P3-level load capacity available during the relevant hour. A shedding scheme is deemed feasible if $S(C)$ does not exceed the magnitude of P3 capacity simultaneously available during the critical hour, ensuring that the shedding decision does not extend to the P1 and P2 levels.

Equipment and Reproducibility

All computations were performed using Python 3, utilizing the pandas and NumPy libraries for time-series processing and Matplotlib for visualization. No imputation or smoothing techniques were applied; consequently, all reported statistical metrics are derived from the original observed values. The procedures for date validation, floor-variable decomposition, and percentile clustering are deterministic and involve no random parameters, ensuring that the results are fully reproducible from the same dataset.

RESULTS AND DISCUSSION

Characteristics of Daily Fluctuations

The average daily load profile in Figure 1 shows that weekend loads remain below weekday loads across all 24 hourly data points. The average weekday load was 1,207.3 kW compared to 1,029.9 kW on weekends, representing a difference of 177.4 kW (−14.7%). A similar pattern is evident in the extreme statistics: the 95th percentile was 1,496.5 kW for weekdays versus 1,244.9

kW for weekends, and the annual peak load was 1,576.2 kW for weekdays versus 1,336.7 kW for weekends. Thus, no weekend load inversion was observed in this study.




Figure 1. Average daily load profiles for weekdays and weekends

From an analytical standpoint, what is far more informative is not the magnitude of the deficit itself, but rather its timing. Between 00:00–02:00 and 22:00–23:00, the difference between the two day categories remains below 1% (ranging from -4.4 kW to -2.8 kW)—a difference that is practically negligible. The deficit emerges at 03:00 (-5.4%), grows steadily to a peak of -390.8 kW (or -26.9%) at 16:00, and subsequently subsides to -2.4% by 21:00. This pattern demonstrates that a hospital's load profile cannot be treated as a monolithic entity that simply fluctuates according to the calendar; instead, it reveals a superposition of two load layers governed by different temporal dynamics: a core layer that remains invariant to the calendar, and a scheduled layer that operates only during specific windows of activity.

The end-use decomposition confirms this distinction while revealing an unexpected composition. The weekend deficit of 177.4 kW comprises interior lighting (-88.7 kW; 50.0% of the total deficit), interior equipment (-44.2 kW; 24.9%), unmetered loads (-40.9 kW; 23.1%), and fans (-10.6 kW; 6.0%). Cooling, conversely, moved in the opposite direction with a slight increase of $+7.2$ kW, rendering it statistically insensitive to the calendar. The technical explanation is both straightforward and significant: the infection control functions performed by the HVAC system—such as air exchange rates, the maintenance of differential pressures in isolation and operating rooms, and the conditioning of sterile storage areas—do not pause for holidays. What diminishes on weekends is the level of outpatient activity, polyclinic operations, administration, and elective support services, rather than the thermal infrastructure or inpatient care operations.

The implications for load-shedding scheme design are direct and—to the best of our knowledge—have not yet been quantitatively defined in the literature on predictive modeling-oriented healthcare facility load profiles (Bagnasco et al., 2015; Cao et al., 2020; Gordillo-Orquera et al., 2018). The shed-able load envelope varies over time, shrinking to near-zero levels during overnight hours. Consequently, a static load-shedding priority table—typical of conventional control panels—is structurally inadequate: at 16:00 on a workday, a substantial load-shedding reserve is available, whereas at 02:00, all remaining loads are critical. A SCADA scheme designed for this profile must be time-aware, rather than merely priority-aware.

It is worth noting that a 14.7% deficit is considered small for a building of comparable size. Office and commercial buildings typically experience much steeper drops in weekend loads as occupancy approaches zero, whereas hospitals maintain continuous operations throughout the week (Soyler & Izgi, 2022). Consequently, the 14.7% figure can be proposed as a quantitative indicator of the degree of operational continuity for a healthcare facility.

Seasonal Distribution and Load Volatility

The monthly distribution box plot in Figure 2 conveys a structural insight often overlooked when analysis stops at average values: intra-month dispersion far exceeds inter-month dispersion. The interquartile range within a single month spans approximately 430 kW, whereas the range of monthly averages is merely 79.7 kW—extending from $1,111.3$ kW in July to $1,191.0$ kW in November, a variation of 7.2% . The ratio between the two is close to $5:1$. The implications for capacity planning

are decisive: seasonality is a second-order effect, while daily cycles are the primary driver of capacity demand. Designs based on monthly billing summaries or monthly averages would fail to capture almost all relevant information.




Figure 2. Monthly electricity load distribution

Regarding the limited month-to-month variation, the cause is attributed to the air conditioning system, as evidenced by the correlation coefficient of $r = 0.845$ between the average monthly total load and the average monthly cooling load. However, the trend runs counter to common expectations: the cooling load is lowest in the middle of the year (469.3 kW in July) and highest during the transitional and cooler months (533.9 kW in November; 531.7 kW in March). The fact that the mid-year period is the hottest for the facility under study is independently confirmed by the gas consumption data, which dropped from 525.0 kW in January to 130.7 kW in July—a pattern characteristic of a climate with cold winters. The mid-year reduction in consumption of 79.7 kW thus comprises a decrease in cooling load of 64.6 kW and a reduction in the non-submetered load of approximately 25 kW, partially offset by an increase in fan load of 24.6 kW (rising from 102.9 kW in January to 127.5 kW in July).

We have chosen not to offer a single explanation for this directional anomaly; at least three hypotheses can be proposed. First, the simultaneous heating and cooling operations in air handling units with high outdoor air fractions, where the cooling coil load during winter is driven by the dehumidification and reheat sequence. Second, the influence of the chiller configuration and staging—along with part-load efficiency—in the prototype model employed. Third, the shift of thermal handling from the chiller side to the air side, consistent with the 23.9% increase in fan load observed mid-year. All three hypotheses require verification of the HVAC system configuration and the weather data underlying the simulation—information not available in the provided dataset. Consequently, this anomaly is noted as both a limitation and a topic for future investigation, rather than an interpreted finding.

Analysis of Base and Peak Load Vulnerability Zones

Figure 3 presents the monthly trajectories of the base, average, and peak loads that form the core of this study's argument. Over the twelve months, the corridor between the 5th and 95th percentiles maintained a nearly constant width, with annual aggregate values of $Q_{0.05} = 880.2$ kW and $Q_{0.95} = 1,484.9$ kW, resulting in a dynamic range of $\Delta P = 604.7$ kW—or 38.4% of the annual peak load of 1,576.2 kW. The annual load factor was recorded at 0.734.




Figure 3. Monthly base and peak load dynamics

Before proceeding further, a conceptual correction must be emphasized. The value of 880.2 kW represents a statistical percentile of the facility's total power distribution; this figure is not a safety threshold for medical equipment—nor should it be interpreted as such—since it is not derived from the power requirements of any specific medical device. The actual, technically sustainable physical limits are defined by two other values: the annual absolute minimum of 782.8 kW and the aggregate floor-level component across all channels—totaling 695.7 kW—which constitutes the P1 level in the Table 1 taxonomy. It is this figure of 695.7 kW—not 880.2 kW—that represents the constant power demand present throughout the entire 8,760-hour observation period; consequently, it serves as the lower limit that must remain untouched by any load-shedding mechanism.

The dynamic range decomposition based on Equation (9) yields the findings that constitute the novelty of this paper, as summarized in Table 2. The conditional mean difference of 651.0 kW between the peak and base regimes comprises interior lighting (198.2 kW; 30.4%), interior equipment (190.6 kW; 29.3%), unmetered loads (178.2 kW; 27.4%), cooling (67.9 kW; 10.4%), and fans (16.1 kW; 2.5%). Cooling—which accounts for 43.3% of annual energy consumption and represents the largest energy component—turns out to contribute only one-tenth of the fluctuation that determines capacity. Its coefficient of variation is 8.6%, compared to 63.4% for lighting, and its correlation with the total load at an hourly resolution is virtually zero ($r = -0.034$). Figure 4 illustrates this visually: the cooling component appears as a nearly flat band across the 24 hours, whereas the lighting component swings from 35 kW to 243 kW. In other words, cooling is an energy load rather than a peak load; it forms the floor, not the roof.

Table 2. Decomposition of the base-to-peak dynamic range by end-use

End use	Basic regime (kW)	Peak regime (kW)	Difference (kW)	Contribution K_i (%)	Energy share (%)
Interior lighting	45.0	243.3	198.2	30.4	11,4
Interior fittings	116.5	307.2	190.6	29.3	17,9
Other loads	143.1	321.1	178.2	27.4	17,7
Cooling system	453.7	521.6	67.9	10.4	43,3
Fan	101.4	117.5	16.1	2.5	9,6
Total	859.7	1,510.7	651.0	100.0	100,0




Figure 4. Decomposition of the daily profile by end-use

Mapping vulnerability zones using Equations (10)–(12) translates these structural findings into design parameters. At a capacity threshold of 1,400 kW, the load exceeds capacity for 2,131 hours (24.3% of the time), with a maximum exceedance magnitude of 176.2 kW and an energy value of 121.8 MWh. At the 1,300 kW threshold, this exposure extends to 2,820 hours (32.2%), with a maximum magnitude of 276.2 kW and an energy value of 378.9 MWh. The feasibility of load shedding is then evaluated against the availability of P3-level capacity: during peak operation, the variable components—lighting, interior equipment, and un-submetered loads—simultaneously total approximately $208.3 + 193.9 + 211.9 \approx 614$ kW. The required load-shedding magnitude for the 1,300 kW scenario thus equates to 45% of the available P3 capacity, allowing for a 17.5% reduction in generator capacity relative to the annual peak load without affecting P1 and P2 levels at all. This quantitative justification forms the basis for developing a SCADA-based load-shedding algorithm—aligned with the tiered medical load classification approach proposed for healthcare facilities on low-

reliability grids (Ibiam et al., 2025)—and allows for integration with power dispatch and energy storage architectures at the microgrid level (Escobar et al., 2024).

Two scientific caveats must accompany this conclusion, and both are substantive. The first concerns the nature of the intervention. An exposure rate of 32.2% of operating time is clearly not characteristic of a rare emergency event. A scheme that operates for one-third of the hours in a year is functionally no longer emergency load shedding, but rather continuous demand limiting. This distinction is not merely a matter of terminology, as the two approaches require different control architectures, hysteresis logic, and operational acceptance criteria.

Secondly, there is the issue of data resolution limits. Generator failure due to tripping is essentially a sub-second transient phenomenon triggered by step loads during power transfer, motor inrush currents, and momentary frequency drops. Hourly resolution data cannot capture such events. Consequently, while this analysis is valid for justifying steady-state capacity settings and formulating load-shedding priorities, it is insufficient to substantiate claims regarding the prevention of generator trips. Addressing the transient regime requires high-resolution data and dynamic stability studies, which we propose as the immediate next step for this research.

CONCLUSION

This study characterizes the hourly electrical load profile of a hospital building—spanning 8,760 observations—and yields four mutually reinforcing quantitative findings. First, weekend loads fall below weekday levels by a margin of 177.4 kW (14.7%)—a notably small deficit compared to similarly sized commercial buildings—with the difference concentrated entirely within specific activity windows: the gap is less than 1% during the 00:00–02:00 and 22:00–23:00 intervals but reaches −390.8 kW (−26.9%) at 16:00, whereas cooling loads remain insensitive to the calendar (+7.2 kW). Second, intra-month load variation (approx. 430 kW) exceeds inter-month variation (79.7 kW) by a ratio of nearly 5:1, demonstrating that seasonality plays only a second-order role in capacity demand. Third, the dynamic range between the base load of 880.2 kW (5th percentile) and the peak load of 1,484.9 kW (95th percentile) spans 604.7 kW—representing 38.4% of the annual peak load of 1,576.2 kW—with a load factor of 0.734, while the critical load present year-round measures 695.7 kW. Fourth—and representing the study's primary novelty—a decomposition of this dynamic range reveals a decoupling between energy share and peak contribution: cooling loads account for 43.3% of annual energy consumption but contribute only 10.4% to the dynamic range, whereas lighting and interior equipment collectively account for 59.7%. As stated in the Introduction, these findings correct the design assumption that identified the HVAC system as the primary driver of variability, shifting intervention priorities toward scheduled loads; subsequent vulnerability zone mapping indicates that reducing generator capacity by 17.5% of the annual peak load is a viable strategy, with a maximum shed magnitude of 276.2 kW—fully encompassed by the approximately 614 kW of deferrable load capacity—without affecting vital or essential loads.

Further development focuses on four agendas addressing the limitations of this study. The first priority is the formulation and testing of a time-aware, SCADA-based automatic load-shedding algorithm; analysis shows that the shed-able load envelope varies by time of day and shrinks to near-zero during the night, rendering static priority tables structurally inadequate. This algorithm must incorporate hysteresis logic and recovery thresholds, given that instances of threshold exceedance account for 32.2% of operating time in the 1,300 kW capacity scenario—functionally transforming the scheme from an emergency intervention into a continuous demand-limiting measure. The second priority involves acquiring high-resolution (second-by-second) data and conducting dynamic stability studies, as generator failure is a transient phenomenon triggered by step-load changes during power transfer and motor inrush currents—dynamics not captured by hourly data. The third priority is validating the model against field measurements from a hospital in a tropical climate using sub-meters to segregate critical branch, essential, and equipment branch panels in accordance with healthcare facility installation standards; this would allow the current proxy-based priority taxonomy to be replaced with an actual classification and enable testing of whether the ki proportions hold true in a climate without a winter season. The fourth priority entails verifying anomalies in monthly cooling load trends against the HVAC system configuration and the weather datasets used in the simulation, as well as extending the analysis to integrate photovoltaics and

battery energy storage by leveraging the persistent base-load characteristics observed overnight (Ali et al., 2023; Leon et al., 2023).

REFERENCES

- Alhurayess, S., & Darwish, M. K. (2012). Analysis of energy management in hospitals. *47th International Universities Power Engineering Conference (UPEC)* (pp. 1–4). IEEE.
- Ali, M., Vasquez, J. C., Guerrero, J. M., Guan, Y., Golestan, S., De La Cruz, J., Koondhar, M. A., & Khan, B. (2023). A comparison of grid-connected local hospital loads with typical backup systems and renewable energy system based ad hoc microgrids for enhancing the resilience of the system. *Energies*, 16(4), 1918. <https://doi.org/10.3390/en16041918>
- Bagnasco, A., Fresi, F., Saviozzi, M., Silvestro, F., & Vinci, A. (2015). Electrical consumption forecasting in hospital facilities: An application case. *Energy and Buildings*, 103, 261–270. <https://doi.org/10.1016/j.enbuild.2015.05.056>
- Bagnasco, A., Saviozzi, M., Silvestro, F., Vinci, A., Grillo, S., & Zennaro, E. (2014). Artificial neural network application to load forecasting in a large hospital facility. *International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)* (pp. 1–6). IEEE.
- Balaras, C. A., Gaglia, A. G., Georgopoulou, E., Mirasgedis, S., Sarafidis, Y., & Lalas, D. P. (2007). European residential buildings and empirical assessment of the Hellenic building stock, energy consumption, emissions and potential energy savings. *Building and Environment*, 42(3), 1298–1314. <https://doi.org/10.1016/j.buildenv.2005.11.001>
- Bertolini, S., Massucco, S., Silvestro, F., Grillo, S., & Giacomini, G. (2013). Short-term load forecasting in hospital systems. In M. Long (Ed.), *World Congress on Medical Physics and Biomedical Engineering: IFMBE Proceedings* (Vol. 39). Springer. https://doi.org/10.1007/978-3-642-29305-4_179
- Cao, L., Li, Y., Zhang, J., Jiang, Y., Han, Y., & Wei, J. (2020). Electrical load prediction of healthcare buildings through single and ensemble learning. *Energy Reports*, 6, 2751–2767. <https://doi.org/10.1016/j.egyr.2020.10.005>
- Deru, M., Field, K., Studer, D., Benne, K., Griffith, B., Torcellini, P., Liu, B., Halverson, M., Winiarski, D., Rosenberg, M., Yazdani, M., Huang, J., & Crawley, D. (2011). *U.S. Department of Energy commercial reference building models of the national building stock* (NREL/TP-5500-46861). National Renewable Energy Laboratory. <https://doi.org/10.2172/1009264>
- Dong, J., Zhu, L., Su, Y., Ma, Y., Liu, Y., Wang, F., Tolbert, L. M., Glass, J., & Bruce, L. (2018). Battery and backup generator sizing for a resilient microgrid under stochastic extreme events. *IET Generation, Transmission & Distribution*, 12(20), 4443–4450. <https://doi.org/10.1049/iet-gtd.2018.5883>
- Escobar, E. D., Betancur, D., & Isaac, I. A. (2024). Optimal power and battery storage dispatch architecture for microgrids: Implementation in a campus microgrid. *Smart Grids and Sustainable Energy*, 9(2), 27.
- Farquharson, D., Jaramillo, P., & Samaras, C. (2018). Sustainability implications of electricity outages in sub-Saharan Africa. *Nature Sustainability*, 1(10), 589–597.
- Gordillo-Orquera, R., Lopez-Ramos, L. M., Muñoz-Romero, S., Iglesias-Casarrubios, P., Arcos-Aviles, D., Marques, A. G., & Rojo-Álvarez, J. L. (2018). Analyzing and forecasting electrical load consumption in healthcare buildings. *Energies*, 11(3), 493. <https://doi.org/10.3390/en11030493>
- Hu, S., Chen, J., & Chuah, Y. (2004). Energy cost and consumption in a large acute hospital. *International Journal of Architectural Science*, 5(1), 11–19.
- Ibiam, N. H., Kahwash, F., & Ahmed, J. (2025). Priority load management for improving supply reliability of critical loads in healthcare facilities under highly unreliable grids. *Energies*, 18(6), 1343. <https://doi.org/10.3390/en18061343>
- Katsanis, J. S., Malahias, G. N., Koustellis, J. D., & Halaris, P. G. (2011). Electrical energy consumptions in hospitals: The case of lifts. *IASTED International Conference on Power and Energy Systems (Europe)*. IASTED.
- Leon, L. M., Romero-Quete, D., Merchan, N., & Cortes, C. A. (2023). Optimal design of PV and hybrid storage based microgrids for healthcare and government facilities connected to highly intermittent utility grids. *Energy Reports*, 9, 5449–5457.
- Morinigo-Sotelo, D., Duque-Perez, O., Garcia-Escudero, L. A., Fernandez-Temprano, M., Fraile-Llorente, P., Riesco-Sanz, M. V., & Zorita-Lamadrid, A. L. (2011). Short-term hourly load forecasting of a hospital using an artificial neural network. *International Conference on Renewable Energies and Power Quality (ICREPQ)* (pp. 441–446).
- Neal, B. (2013). Spec'ing hospital electrical distribution system. *Consulting-Specifying Engineer*, 50, 22–27.
- Palensky, P., & Dietrich, D. (2011). Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE Transactions on Industrial Informatics*, 7(3), 381–388. <https://doi.org/10.1109/TII.2011.2158841>
- Saki, H., Zangeneh, A., & Aghaei, J. (2025). Distributed minimum spanning tree approach for critical load restoration using microgrid formation in resilient distribution systems. *Electric Power Systems Research*, 239, 111186.

Soyler, I., & Izgi, E. (2022). Electricity demand forecasting of hospital buildings in Istanbul. *Sustainability*, 14(13), 8187. <https://doi.org/10.3390/su14138187>