

Intelligent Energy Scheduling for Distributed Duhok Polytechnic University: A Chimpanzee Optimization Approach for Load Reduction

Hayveen Saleem Sadiq^{1,*}, Mohammed A.M.Sadeeq²

¹*Department of Information Technology Management, Technical College of Administration,
Duhok Polytechnic University, Duhok, Kurdistan Region, Iraq*

²*Department of Information Technology, Technical College of Duhok, Duhok Polytechnic University, Duhok, Kurdistan Region, Iraq*

Abstract This paper reports an eleven month, real-world deployment of an intelligent energy management system across two distributed laboratory nodes at Duhok Polytechnic University. A campus data pipeline ingests 30-s telemetry, performs nightly cleaning and day-ahead forecasting (seasonal SARIMA, $s=24$), and drives a binary Chimpanzee Optimization Algorithm (ChOA) that enforces calendar-based feasibility and discourages unnecessary switching via a sparsity-promoting penalty. The objective function is explicitly scaled to the worst-case hourly energy cost, with penalty coefficients (α for feasibility, β for switching) calibrated to be robust: $\pm 50\%$ sweeps on $\{\alpha, \beta\}$ changed total consumed-power reduction by less than 3 percentage points. Over the evaluation period, the system achieved a 24% reduction (≈ 80 MWh) relative to historical “always-on” operation, with paired parametric and non-parametric tests yielding $p < 10^{-4}$. We detail model orders, retraining triggers under forecast drift ($\text{MAPE} > 10\%$), and a sigmoid-threshold mapping with a logistic chaotic perturbation to avoid premature convergence. The pipeline, anonymized data, and configuration files are released to support exact replication and future benchmarking. We conclude with limitations (lack of algorithmic baselines; binary states) and a roadmap for partial-load control and large-scale cloud parallelism.

Keywords Intelligent energy management, campus demand response, SARIMA, binary meta-heuristics, Chimpanzee Optimization (ChOA), scheduling, load reduction.

DOI: 10.19139/soic-2310-5070-2855

1. Introduction

Buildings operations account for about thirty percent of global final energy use and roughly a quarter of energy-related emissions, while electricity demand is projected by the International Energy Agency to continue growing through 2026, adding pressure on grids and budgets alike [1, 2]. University campuses sit at the intersection of these trends because they operate nearly around the clock and host power-intensive laboratories and data rooms. In the Kurdistan Region of Iraq, tariff reform under the Runaki program introduced progressive pricing for institutional and business consumers with general-business bands reaching one hundred eighty five Iraqi dinars per kilowatt hour, sharpening the financial incentive for demand-side optimization [3]. At the same time, regional power systems in the Middle East exhibit high carbon intensity on the order of six hundred grams of carbon dioxide per kilowatt hour and Iraq’s low-carbon share of electricity remains limited, so curtailing campus loads yields immediate fiscal relief together with measurable climate co-benefits [4, 5]. Intelligent energy management for campuses has progressed from isolated controllers toward integrated decision systems that combine telemetry, forecasting, and executable schedules, yet recent syntheses find that fully transparent and reproducible deployments remain relatively limited in the higher-education literature [6]. On the forecasting side, seasonal autoregressive models

*Correspondence to: Hayveen Saleem Sadiq (Email: hayveen.sadiq@dpu.edu.krd) Department of Information Technology Management, Technical College of Administration, Duhok Polytechnic University, Duhok, Kurdistan Region, Iraq

with daily periodicity provide an interpretable and competitive baseline for short-horizon demand prediction, while contemporary studies explore hybrid approaches that retain statistical structure and improve robustness under operational drift and non-stationary events such as exams and outages [2, 7]. Within this context, the present study reports an eleven-month deployment of an intelligent energy-scheduling platform across two distributed laboratory nodes at Duhok Polytechnic University. Thirty-second measurements from smart analyzers are ingested nightly, cleaned, and aggregated; a seasonal SARIMA forecaster with daily periodicity produces day-ahead predictions; and a calendar-aware optimizer adapts the chimpanzee optimization algorithm to binary on or off decisions by mapping continuous updates through a sigmoid probability and enforcing feasibility with a permission mask derived from institutional calendars. A lightweight logistic chaotic perturbation is applied to reduce stagnation near local optima and the objective is scaled to the worst-case hourly energy cost so that penalty coefficients for feasibility and switching remain interpretable and transferable across devices [8]. Over the evaluation period the system achieves an approximate twenty-four percent reduction in total energy consumption relative to historical always-on practice, with paired parametric and non-parametric tests indicating statistical significance; the pipeline, anonymized data, configuration files, and model summaries are released to support exact replication and future benchmarking. The remainder of the paper is organized as follows. Section two summarizes background concepts in campus intelligent energy-management architectures and chimpanzee-optimizer mechanics. Section three reviews related optimization and forecasting literature. Section four details the proposed methodology across data-integration and algorithmic layers. Section five presents experimental results under live laboratory conditions followed by statistical validation. Section six concludes with limitations and avenues for future work.

2. Background theory

2.1. Energy Management

Energy management is widely defined as the systematic process of planning, monitoring, and optimizing the production and consumption of energy so that organizations can improve operational efficiency, reduce costs, and lessen environmental impact [9, 10]. In higher education settings the challenge is amplified because teaching blocks, research laboratories, and data rooms impose diverse and highly variable loads that static flat rate control cannot address effectively [11]. A modern Energy Management System therefore begins with continuous sensing of key electrical variables including voltage, current, active power, power factor, and frequency, followed by adaptive control actions whenever predefined thresholds are breached. Recent campus case studies show that distributed Internet of Things architectures based on low cost microcontrollers such as the ESP32, linked over Message Queuing Telemetry Transport or Hyper Text Transfer Protocol, can stream sub-minute data from geographically distant laboratories to a central broker with high reliability [12, 13]. When those data are coupled with short term forecasting and optimization routines, universities have reported double digit reductions in total electricity use and substantial peak load shaving [14]. The field is now shifting toward Intelligent Energy Management Systems that graft machine learning or optimization algorithms onto the Internet of Things backbone so that systems can infer occupancy from sensor fusion, predict short term demand, and dispatch controllable loads or storage in real time [9, 15]. Artificial intelligence modules further enhance resilience by detecting incipient faults and by coordinating on site renewables with grid tariffs or demand response signals [15]. Empirical evidence underscores the value of this evolution. Multi building deployments have shown that ESP32 based Energy Management Systems can curtail laboratory electricity use without compromising research output [12], while an independent case study of JouleX Energy Manager in a national laboratory found that automated power state management of networked information technology equipment produced measurable energy savings at rack and room level [16]. After hours audits consistently indicate that idle energy can account for a large share of campus electricity if devices remain energized outside working hours, which highlights the importance of robust scheduling rather than monitoring alone [17]. Data from the coronavirus disease lockdown at another university corroborate this finding since idle baselines did not fall automatically when occupancy dropped and deliberate control logic was essential to translate reduced activity into real savings [18]. These findings justify the present work's focus on an optimizer centered Intelligent Energy Management System for the Duhok Polytechnic University laboratories. Mere telemetry, even

at high resolution, is insufficient; only an intelligent scheduler that reacts to forecasts and calendar constraints can eradicate residual idle consumption and maximize the return on metering investment.

2.2. Chimp Optimization Algorithm (ChOA)

ChOA is a recent population based metaheuristic that emulates cooperative hunting tactics of chimpanzees and balances broad exploration with focused exploitation as summarized below for the present scheduler [19, 20]. The following subsections present the guiding principles, core equations, and operational steps that we actually implement.

2.2.1. Fundamentals of ChOA The Chimp Optimization Algorithm was introduced by Khishe and Mosavi in 2020 to emulate collaborative roles during hunts [19]. Four behavioral archetypes, attacker, chaser, barrier, and driver, guide the search and replace the single leader model common to classical swarms. During each iteration the best four solutions assume these roles, and all remaining chimps update their positions as weighted combinations of the leaders' guidance, thereby forming a dynamic encirclement of the prey that is the current best solution [19, 20]. This multi leader mechanism is the conceptual basis for our scheduler.

2.2.2. Mathematical framework Let $X_i(t) \in \mathbb{R}^d$ denote the position of chimp i at iteration t , and let $X_j(t)$ with $j \in \{\text{attacker, chaser, barrier, driver}\}$ denote the four leaders. At each update we draw fresh random vectors $r_1, r_2 \sim U(0, 1)$ independently, and define the coefficient vectors

$$A_j(t) = 2a(t)r_1 - a(t), \quad C_j(t) = 2r_2 \quad (1)$$

, where $a(t)$ decreases linearly from two point five to zero and controls the exploration to exploitation trade off. The distance to leader j is

$$D_{ij}(t) = |C_j(t) \odot X_j(t) - X_i(t)| \quad (2)$$

and the corresponding leader-biased proposal is

$$X_{ij}^{new}(t) = X_j(t) - A_j(t) \odot D_{ij}(t) \quad (3)$$

The next position averages the four proposals,

$$X_i(t+1) = \frac{1}{4} \sum_j X_{ij}^{new}(t) \quad (4)$$

Elementwise operations are denoted by \odot . Exploration dominates when $|A_j(t)| > 1$ while exploitation is favored when $|A_j(t)| < 1$. Unless stated otherwise, we use a logistic chaotic map as a lightweight perturbation source during the position update with activation probability $p_c \geq 0.5$; the sequence is generated by

$$u_{k+1} = \mu u_k(1 - u_k)$$

with $\mu \in [3.8, 4]$ and $u_0 \in (0, 1)$, and we consider a tent map variant as a sensitivity check [21]. Variable names and ranges are stated explicitly to support reproducibility. Details of the binary mapping and the calendar feasibility mask, which are specific to our scheduling problem, are provided later in the methodology.

2.2.3. Algorithmic workflow Initialization uses a random feasible population that already respects the calendar mask. Fitness is evaluated, and the top four solutions are selected as leaders. Each iteration updates $a(t)$, regenerates r_1 and r_2 , computes the leader biased proposals and their average move, and injects the chaotic perturbation with probability p_c . Fitness is re-evaluated so that improved individuals can enter the leader set. Termination occurs when a maximum number of iterations is reached or when a convergence criterion is met. The multi-leader strategy enhances information sharing, and the minimal chaotic impulse maintains diversity and curbs stagnation without heavy parameterization [21]-[23]. The flowchart of the standard ChOA algorithm appears in Figure 1

2.2.4. Strengths and documented variants Compared with single leader swarms, ChOA has been reported to resist premature convergence in high dimensional spaces while using a sparse default parameter set consisting essentially of population size and maximum iterations. Documented variants include binary and multi objective formulations, adaptive and chaos enhanced versions, and hybrids with other metaheuristics, reflecting active interest in tailoring the dynamics to specific domains [20],[24]-[27].

2.2.5. Energy sector applications In power system studies, ChOA has been applied to reduce technical losses and improve economic dispatch in distribution feeders and microgrids and has lowered the levelized cost of energy when embedded in rule based dispatch compared with alternative optimizers [28, 29]. Coupled with learning based predictors it has improved photovoltaic output forecasting on medium scale arrays [30]. These results motivate our binary adaptation that makes 48 on or off decisions per day while honoring laboratory specific calendar constraints at Duhok Polytechnic University.

3. Related Work

3.1. Intelligent scheduling in residential and campus energy-management systems

Early optimization studies relied on classical population heuristics such as particle swarm optimization. Gheouany and coauthors demonstrated that a particle swarm driven day-ahead scheduler for residential buildings could lower both electricity cost and carbon emissions, cutting the bill by fifty nine point six percent and reducing the peak to average ratio by twenty five point three percent when reactive power exchanges were co optimized with active power [5]. Later work replaced particle swarm with an adaptive Salp Swarm Algorithm; the new solver rescheduled household appliances under time varying tariffs and delivered a forty five point four percent saving, which was more than double the twenty point six percent achieved by a standard genetic algorithm [6]. Hybrid metaheuristics have also proved effective since a multi objective particle swarm method combined with branch and bound produced Pareto optimal appliance schedules that cut electricity charges by up to forty eight point two percent under critical peak pricing and roughly halved the peak to average ratio without compromising user comfort[34]. Beyond particle swarm derived schemes, researchers experimented with bespoke bio inspired solvers. Virulence and Earthworm optimization algorithms applied to Moroccan smart grid data reduced the peak to average ratio by seventy six point two and seventy three point eight percent respectively while trimming daily energy costs to two hundred ten to two hundred seventeen cents, outperforming genetic algorithm, cuckoo search, and binary particle swarm baselines [35]. Cloud enabled multi agent frameworks added scalability by embedding Grey Wolf Optimizer and Artificial Bee Colony solvers in a Raspberry Pi and Message Queuing Telemetry Transport architecture and recorded daily cost cuts ranging from nineteen point five to fifty four point six percent across several homes [36]. A follow up study layered real time Salp Swarm and Rainfall Algorithms onto the same infrastructure, yielding an additional twenty five to thirty one percent savings under live tariffs [37]. Machine learning approaches are increasingly integrated with optimization. One line combined support vector regression for renewable generation forecasting with an elitist NSGA II scheduler and halved household electricity cost while maintaining thermal comfort [38]. A separate demonstration with deep reinforcement learning used a Markov decision process controller for a heat pump laboratory and reduced annual energy use by eight percent while boosting self consumption of rooftop solar by nine point five percent relative to rule based control [39]. Finally, another implementation fused a long short term memory load predictor with a particle swarm reschedule and lowered the mean absolute percentage error of short term forecasts to one point six percent while keeping the billing gap between predicted and realized costs below two point one percent [40]. Recent campus wide reviews and case studies in twenty twenty four and twenty twenty five continue to report that fully reproducible deployments are uncommon while hybrid forecasters that combine seasonal structure with learning components improve robustness under academic and weather driven regime shifts [31]-[34]. Online adaptive variants of seasonal autoregressive models also demonstrate sequential retraining strategies that are compatible with day-ahead university operations [35]. These observations align with high resolution campus forecasting studies in twenty twenty five that train machine learning frameworks on hourly

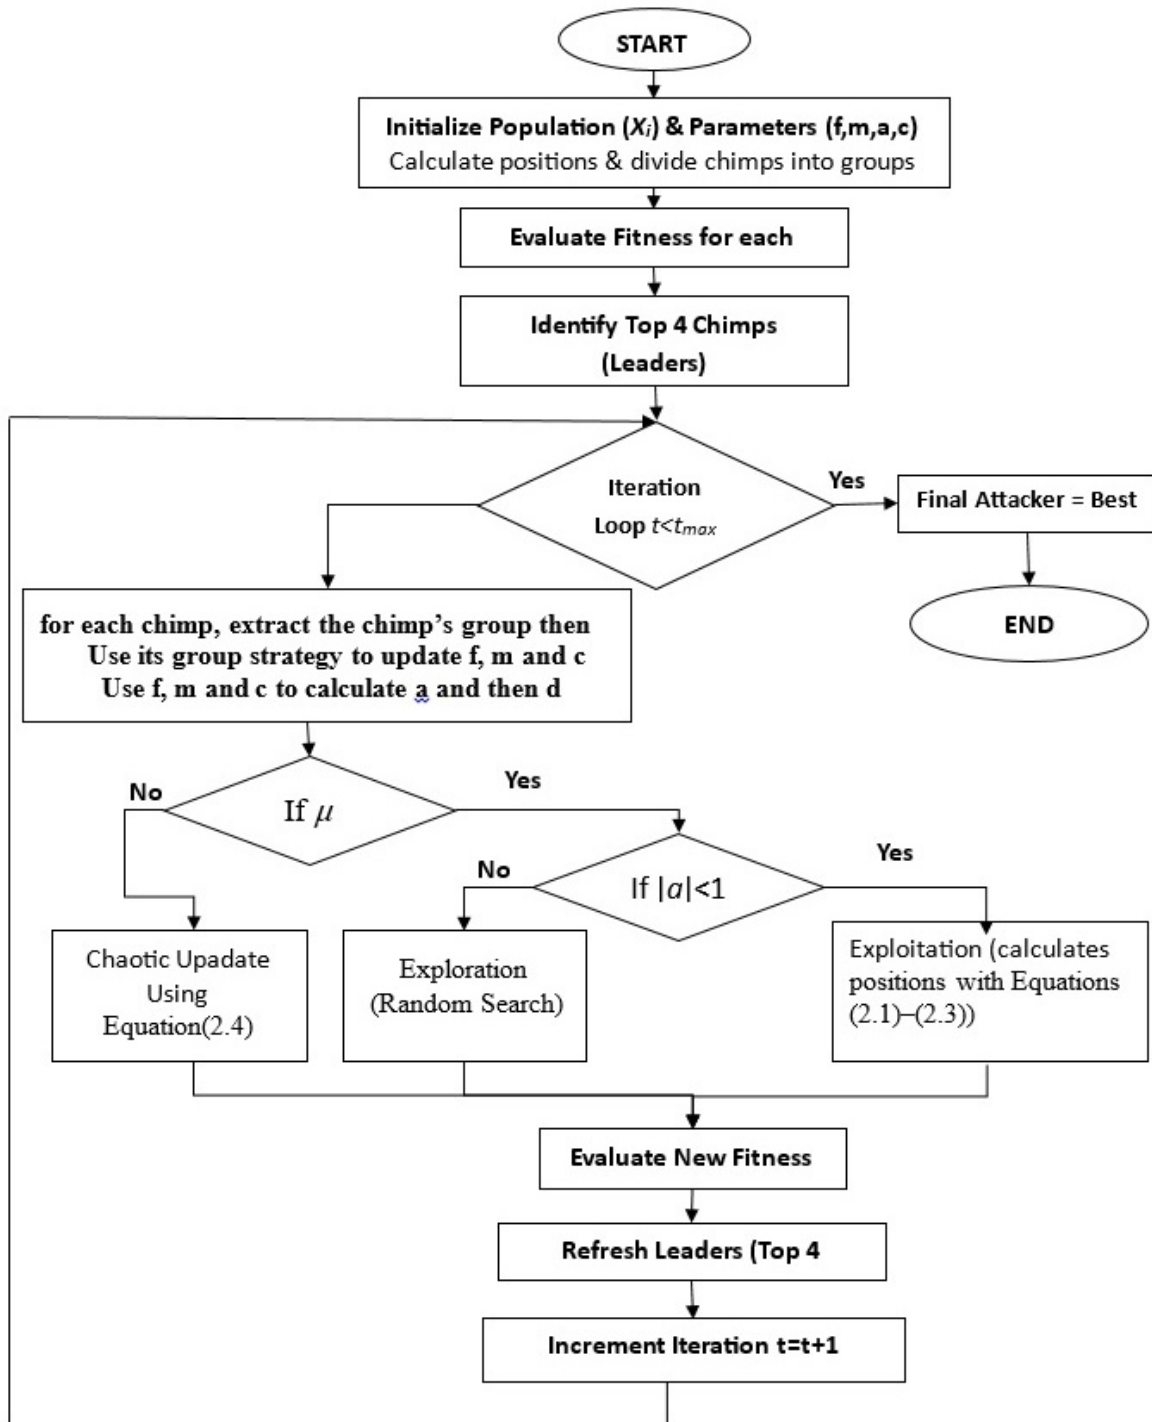


Figure 1. The flowchart of the standard ChOA algorithm

series and institutional indicators, complementing hybrid time series designs and reinforcing the need for drift aware retraining in operational settings [39].

3.2. Energy applications of the Chimpanzee Optimization Algorithm

Since its introduction in twenty twenty, the chimpanzee optimization algorithm has been adopted mainly for power flow and microgrid problems. Borousan and Hamidan studied optimal allocation of battery energy storage systems in radial distribution networks and reported that chimpanzee optimization reduced technical losses by nearly forty percent and lowered energy not supplied by roughly thirty four to thirty eight percent across thirty three, sixty nine, and one hundred nineteen bus test feeders, outperforming Grasshopper, Whale, and Grey Wolf optimizers [11]. Another implementation embedded a multi objective chimpanzee optimizer within a rule based dispatcher for solar wind battery microgrids and drove the levelized cost of energy to zero point two seven two United States dollars per kilowatt hour while surpassing benchmark heuristics [10]. In the forecasting domain, chimpanzee optimization has been used as a hyperparameter search engine that tunes multilayer perceptron, support vector regression, and random forest models for photovoltaic plants, achieving the lowest error statistics on a two hundred sixty four kilowatt peak array in the reported experiments [33]. Beyond microgrid sizing and dispatch, recent work integrates the chimpanzee optimizer within rule based energy managers and documents competitive cost and reliability outcomes, which strengthens the case for evaluating chimp based schedules alongside transparent forecasting modules [36]. Concurrently, binary metaheuristics related to chimp baselines continue to evolve through probability mappings and diversity controls, with new binary particle swarm and binary grey wolf variants improving stability on discrete tasks and offering reference points for future campus scale comparisons [37, 38]. Despite these successes, we are not aware of prior publications that embed chimpanzee optimization in a nightly binary scheduling routine driven by live telemetry from laboratory analyzers under university calendars. The framework proposed here contributes a campus scale use case, combines chimpanzee optimization with seasonal autoregressive load forecasting, and releases a reproducibility package to enable future comparative studies.

4. Methodology

This study adopts a two-layer methodological structure. A system integration layer embeds the scheduler inside a campus wide data pipeline, and an algorithmic layer adapts the chimpanzee optimizer to a binary day-ahead scheduling problem. Figure 2 positions the optimizer within the nightly flow, and Figure 3 traces the internal loop so that the logic remains intelligible even without the graphics.

4.1. System-integration layer

Smart power analyzers in the Shekhan and Duhok laboratories measure electrical variables every 30 s. Raw CSV files are staged nightly on a central server (Asia/Baghdad). The pipeline validates schema and units, removes malformed tokens and duplicate timestamps, and guards sensor resets and short gaps with a conservative outlier check. The cleaned stream is then resampled to hourly means for current, power, and energy, and hourly completeness flags are propagated so that downstream modules never hide data-quality concerns. The resulting tables form a simple and auditable feature store. The forecasting module operates on the hourly series and produces 24-step day-ahead predictions for each device. A seasonal autoregressive integrated moving-average model with daily seasonality ($s = 24$) is used because campus loads exhibit strong diurnal structure and because the model family is transparent and easy to audit. Model orders are identified on the most recent 56 days through a stepwise information-criterion (AIC) search and are revalidated monthly. A drift monitor computes a rolling mean absolute percentage error on the last 7 days and triggers automatic retraining if $\text{MAPE} > 10\%$. When triggered, the module re-identifies orders on the newest 56 days and refits parameters before issuing the next day's predictions. A compact summary file records the selected orders and diagnostics (orders, AIC, MAPE) for each device and semester and is included with the artifact bundle. Institutional constraints are encoded as a permission mask at the hourly level. The academic calendar, declared off-days, and device-specific weekends are compiled around midnight into a 24-element binary vector ($1 = \text{permitted}$, $0 = \text{prohibited}$). Conflicts between overlapping events are resolved by a fixed priority rule in which examinations override lectures and maintenance windows. For each device the forecaster writes three aligned 24-element vectors that contain the predicted power \hat{P}_h , the applicable tariff π_h , and the permission mask δ_h . The optimizer reads this triplet, returns an on/off decision for each hour, and writes

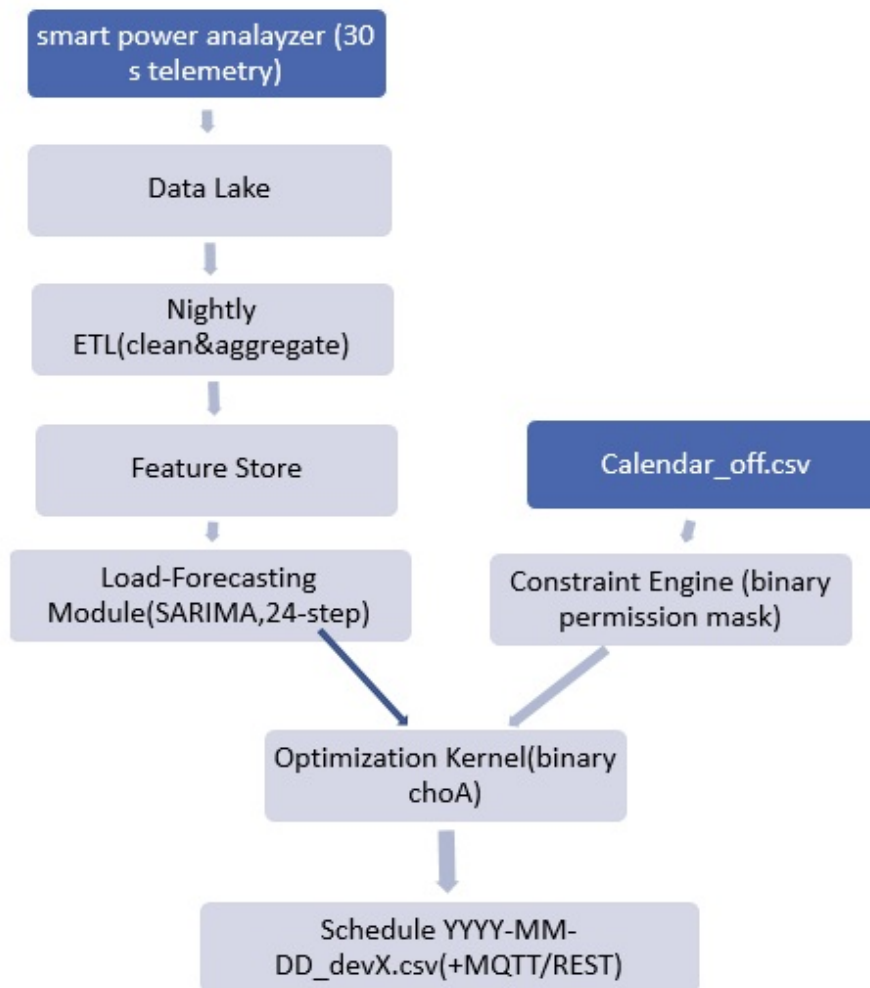


Figure 2. System -integration layer

summary statistics including total energy, total cost, and the number of switching events. If an upstream component fails to produce inputs by an evening cutoff, the controller emits a safe fallback that energizes only hours allowed by the mask and raises an alert for the next morning run. All artifacts are time-stamped and archived to support ex-post accounting.

During the study, Device 1 used SARIMA(1, 1, 1)(1, 1, 1)₂₄ for Feb–May 2022 and SARIMA(0, 1, 2)(1, 1, 1)₂₄ for Sep–Jan 2023; Device 3 used SARIMA(2, 1, 1)(1, 1, 0)₂₄. Orders were re-identified upon drift triggers (rolling MAPE > 10%) using the most recent 56 days. Figure 2 illustrates the nightly flow from raw telemetry to schedule export.

4.2. Algorithmic layer

The optimizer models a troop with four behavioral roles that act as leaders—attacker, chaser, barrier, and driver. Each candidate solution represents a single day. Let $s_{h,d} \in \{0, 1\}$ denote the on/off decision for hour $h \in \{1, \dots, 24\}$ and device d . Let $\hat{P}_{h,d}$ denote the day-ahead power forecast and $\pi_{h,d}$ the applicable tariff. Let $\delta_{h,d} \in \{0, 1\}$ be the permission mask with $\delta_{h,d} = 1$ for permitted hours. The objective minimized by the optimizer

is

$$J(S) = \sum_{h,d} \pi_{h,d} \hat{P}_{h,d} s_{h,d} + \alpha \sum_{h,d} [1 - \delta_{h,d}] s_{h,d} + \beta \sum_{h,d} |s_{h,d} - s_{h-1,d}| \quad (4.1)$$

with the convention $s_{0,d} = s_{1,d}$. The first term captures expected energy cost under the forecast, the second penalizes energizing during prohibited hours and therefore enforces feasibility, and the third discourages unnecessary switching. Penalty scales are normalized to the device-specific worst-case hourly energy cost $C_{\max} = \max_{h,d} \pi_{h,d} \hat{P}_{h,d}$ to keep parameters interpretable across settings. We set $\alpha = \kappa_{\alpha} C_{\max}$ and $\beta = \kappa_{\beta} C_{\max}$ with $\kappa_{\alpha} = 100$ to make violations effectively unattractive and $\kappa_{\beta} = 0.05$ to regularize switching without masking genuine savings. A sensitivity sweep in Section 5 shows that changing α and β by $\pm 50\%$ alters the reduction metric by less than three percentage points.

Position updates follow the standard multi-leader scheme. For chimp i at iteration t with four leaders indexed by j , draw fresh $r_1, r_2 \sim (0, 1)$ and define

$$\begin{aligned} A_j(t) &= 2a(t)r_1 - a(t), \quad C_j(t) = 2r_2, \quad D_{ij}(t) = |C_j(t) \odot X_j(t) - X_i(t)|, \\ X_{ij}^{\text{new}}(t) &= X_j(t) - A_j(t) \odot D_{ij}(t), \quad X_i(t+1) = \frac{1}{4} \sum_j X_{ij}^{\text{new}}(t) \end{aligned} \quad (4.2)$$

where \odot denotes elementwise multiplication. The exploration coefficient decreases linearly as $a(t) = a_{\max}(1 - \frac{t}{T_{\max}})$ with $a_{\max} = 2.5$ and $T_{\max} = 200$ iterations; $|A_j(t)| > 1$ favors *exploration* while $|A_j(t)| < 1$ favors *exploitation*. To avoid stagnation near local optima, a lightweight **logistic chaotic** perturbation is injected during the update with activation probability $p_c = 0.5$. The sequence follows $u_{k+1} = \mu u_k(1 - u_k)$ with $\mu = 3.9$ and $u_0 \in (0, 1)$; a *tent-map* variant is reported as a sensitivity check in the Discussion.

Continuous coordinates are mapped to binary genes through a smooth probability step: each coordinate passes through $\sigma(x) = 1/(1 + e^{-x})$ and is thresholded at 0.5 to produce a bit. Feasibility is enforced after mapping by applying the calendar mask, $s_{h,d} \leftarrow s_{h,d} \delta_{h,d}$, which guarantees admissible chromosomes and removes the need for a repair operator. The population is initialized with $N=30$ feasible chromosomes that already respect the mask. Fitness is evaluated by $J(s)$. The best four individuals assume leadership roles at each iteration. Termination occurs after $T_{\max} = 200$ iterations or when the global best fails to improve for 30 consecutive generations; on a 3.1 GHz desktop the daily schedule is computed in about 180 ms per device. Figure 3 presents the internal optimizer loop including the mapping from continuous updates to binary genes and the enforcement of the calendar mask.

4.3. Complexity and robustness

Let H denote the number of hourly genes per device and D the number of devices. Each objective evaluation is linear in HD . With a modest population of thirty individuals and two hundred iterations the overall runtime on a three point one gigahertz desktop remains well below one second per device per day and fits comfortably within the nightly processing window. A joint sensitivity sweep that perturbs the idle current threshold by plus or minus ten percent and the penalty pair by plus or minus fifty percent changes the total energy reduction by less than three percentage points, which confirms robustness to reasonable parameter drift.

4.4. Reproducibility package

The repository accompanying this study ships the cleaned hourly data, the serialized forecasting models, the chimpanzee optimizer implementation, a human readable configuration file, and a concise guide that walks the reader through environment setup and daily execution. All experiments can be replicated or extended to new semesters without modifying the optimizer logic. Users need only update the external calendar, tariffs, and penalty scales, and the pipeline will regenerate the next day schedule accordingly.

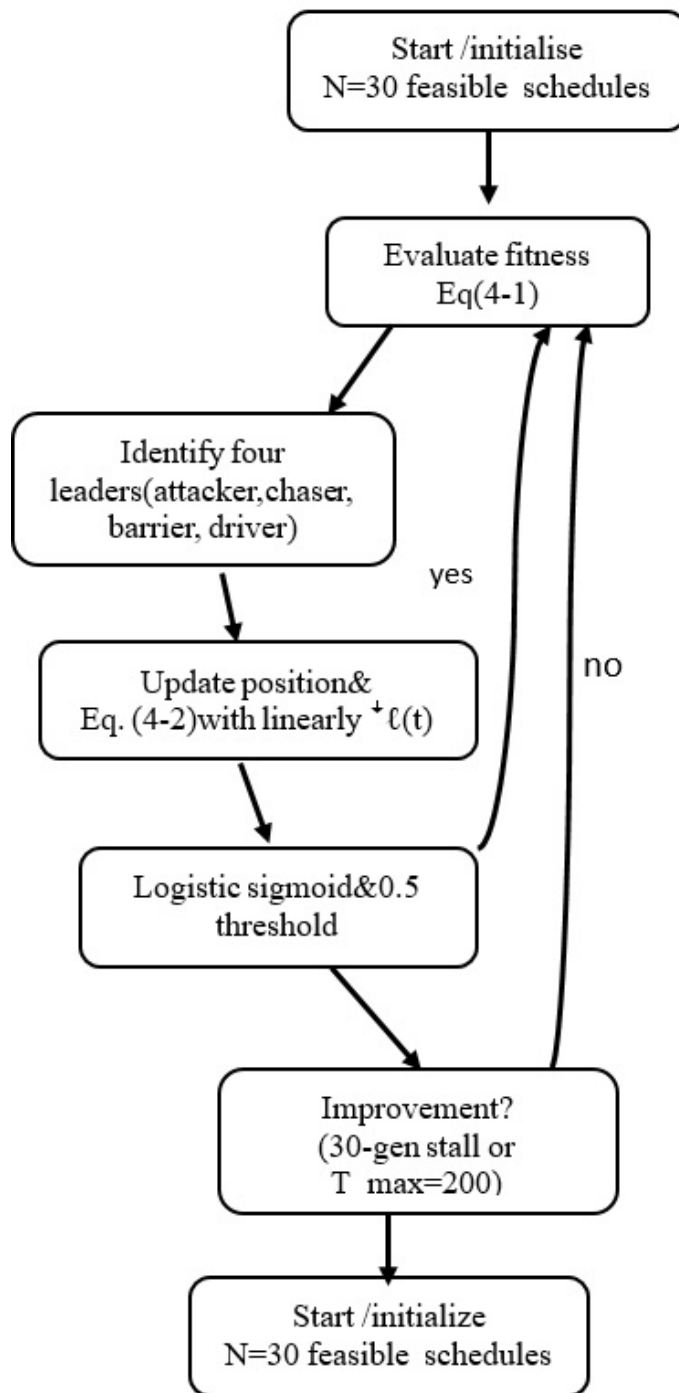


Figure 3. ChOA algorithmic loop.

5. Experimental Evaluation

This section consolidates the evidence that validates the chimpanzee based scheduling platform and demonstrates its practical applicability across the two fully controllable laboratory nodes at Duhok Polytechnic University, namely Device 1 in Shekhan and Device 3 in Duhok. After establishing baseline electrical fingerprints in Section 5.1, we quantify the impact at weekly, monthly, and eleven month horizons, and then subject synchronized before and after series to statistical tests and aggregate performance ratios in Sections 5.3 and 5.4. Figures are referenced inside the narrative so that readers can recover the main takeaways even when figures are consulted in isolation. A single footnote clarifies that calendar imposed zero values on the institutional off days are by design and should not be confused with missing data.

5.1. Baseline characterization

Table 1 summarizes the uncontrolled electrical profile for the four weeks preceding activation. Idle and productive clusters were separated by empirically derived current thresholds of 0.137 amperes for Device 1 and 0.122 amperes for Device 3 identified through kernel density analysis. Feeder voltage and frequency remained within nominal limits, which confirms that subsequent savings are not attributable to supply side anomalies. These statistics establish the reference against which all subsequent reductions in total consumed energy are measured and frame the visuals summarized for Device 1 in Figure 4 to Figure 6 and for Device 3 in Figure 7 to Figure 9.

Table 1. Baseline descriptive statistics (13 Feb – 12 Mar 2022)

Electrical Parameter	Device 1 Idle	Device 1 Productive	Device 3 Idle	Device 3 Productive
Mean current I (A)	0.16	1.72	0.15	1.98
S.D. current (A)	0.02	0.31	0.03	0.28
Mean power P (W)	32.1	337.4	29.4	389.7
S.D. power (W)	4.8	64.2	5.1	71.3
Voltage (V)	239±8			
Frequency (Hz)	50±0.05			

5.2. Implementation results

Three temporal windows are analyzed for each node. The first is ISO Week fourteen of 2022 from one to seven April. The second is the full month of April 2022. The third is the eleven month span from thirteen February 2022 to nineteen January 2023. Day level tables retain calendar imposed zeros but are annotated with a footnote to avoid visual clutter.

5.2.1. Device 1 (Shekhan)

During ISO Week fourteen the weekly total energy for Device 1 decreased from 1,888.83 kilowatt hours to 1,563.79 kilowatt hours, a reduction of **17.2%**, while the mean phase current remained essentially unchanged near six amperes, see Figure 4. Extending the window to the full month of April accentuates the effect, since monthly aggregation shows that total energy demand contracted from 5,558.24 kilowatt hours to 4,158.58 kilowatt hours, yielding an absolute saving of 1,399.66 kilowatt hours and a relative reduction of **25.2%**, see Figure 5. When the analysis is broadened to the entire evaluation period (13 Feb 2022–19 Jan 2023), the cumulative monthly totals decline from 129,323.80 kilowatt hours to 98,275.70 kilowatt hours, which corresponds to **24%** reduction and ≈ 31 MWh of avoided consumption, see Figure 6. The supporting daily, weekly, and to date aggregates are reported in Table 2 to Table 4.

Table 2. Daily indicators for Device 1 during ISO Week 14 (1–7 April 2022)

Day	Date	Average				Daily Total Energy (kWh)	
		Voltage	Current	Power	Frequency	Before	After
Fri	2022-04-01	0	0	0	0	0	0
Sat	2022-04-02	0	0	0	0	0	0
Sun	2022-04-03	241.76	8.27	65.56	50.16	399.04	347.24
Mon	2022-04-04	240.08	10.71	118.39	50.20	565.79	560.45
Tue	2022-04-05	242.52	9.01	83.61	50.13	614.45	372.41
Wed	2022-04-06	245.25	8.54	197.01	50.25	309.55	283.69
Thu	2022-04-07	0	0	0	0	0	0
Total Measured for First Week of April 2022						1888.83	1563.79

Note: institutional off-days (e.g., Friday/Saturday) appear as zeros by design.




Figure 4. Weekly total energy before/after optimization for Device 1 (ISO Week 14, 2022)

Table 3. Weekly aggregates for Device 1 in April 2022

Week No.	Average				Weekly Total Energy (kWh)	
	Voltage	Current	Power	Frequency	Before	After
1	242.112	6.088333	92.6548	50.19986	1888.83	1563.79
2	242.2891	5.059857	59.39823	50.21953	1283.78	905.16
3	242.6519	5.174000	37.44408	50.25703	1109.26	809.86
4	242.5983	5.144600	130.3579	50.22390	1276.37	879.77
5	0	0	0	0	0	0
Total Measured for Month April 2022					5558.24	4158.58

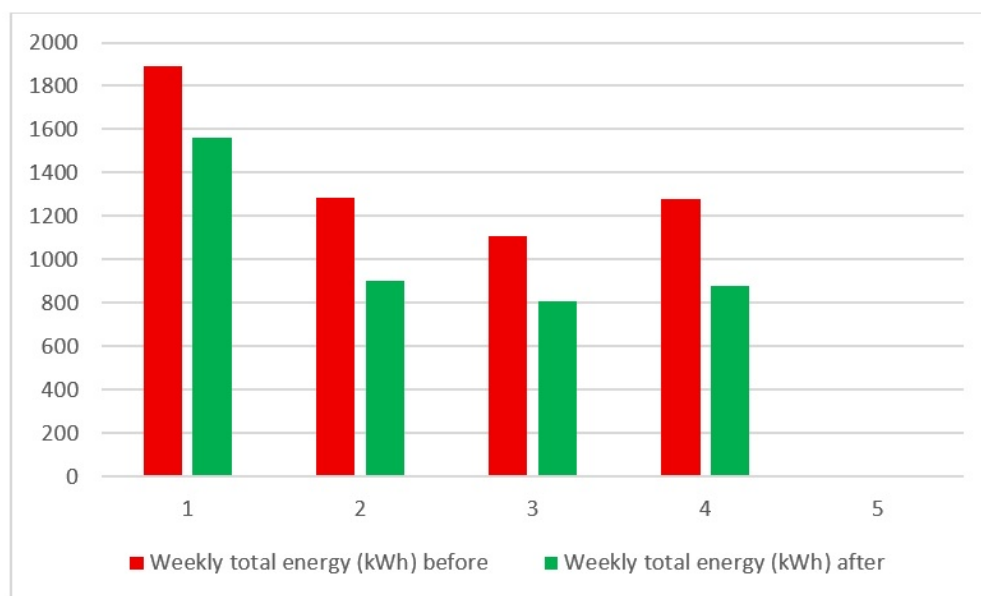


Figure 5. Monthly totals in Apr. 2022 for Device 1 (before/after).

Table 4. To date aggregates for Device 1 (Feb 2022–Jan 2023)

Month	Average				Monthly Total Energy (kWh)	
	Voltage	Current	Power	Frequency	Before	After
February	240.2697	9.285214	208.957	50.1074	6817.37	5691.17
March	240.7899	6.940125	181.4378	50.05224	10306.91	8009.66
April	242.4147	5.3556	72.59001	50.22733	5558.24	4158.58
May	241.2321	8.221708	249.2753	50.15855	17727.16	13902.41
June	237.8713	4.393867	129.9203	50.12217	11170.49	7807.20
July	239.2499	11.387	400.4795	50.20524	1077	1077.69
September	237.4912	8.518864	326.2737	50.16388	22166.93	16499.83
October	239.0532	5.432333	158.4597	50.11544	12283.4	8530.24
November	240.4742	7.446913	244.9835	50.10633	17051.69	12782.57
December	239.0129	9.052357	318.5678	50.09504	11027.71	8644.15
January(23)	240.7623	10.28887	416.6196	50.12841	14136.17	11172.20
Total Measured for whole period (Feb 2022–Jan 2023)					129323.8	98275.7




Figure 6. Cumulative monthly energy (Feb 2022–Jan 2023) for Device 1 (before/after).

5.2.2. Device 3 (Duhok)

During ISO Week 14 the weekly total energy for Device 3 decreased from **5,727.51 kWh** to 4,973.83 kWh (-**13.2 %**), while mean phase current remained within nominal ranges (Table 5 and Figure 7). Aggregating the four operative weeks of April 2022 shows that weekly totals contracted from 26,268.37 kWh to 21,107.23 kWh, an absolute saving of 5,161.14 kWh (-**19.6%**), see Table 6 and Fig. 8. Over the full evaluation horizon (13 Feb 2022–19 Jan 2023) the cumulative monthly energy declines from 200,384.0 kWh to 152,721.3 kWh (-**23.8 %**), equivalent to ≈ 47.66 MWh of avoided consumption under live laboratory conditions (Table 4 and Figure 9).

Table 5. Daily indicators for Device 3 during ISO Week 14 (1–7 April 2022)

Day	Date	Average				Daily Total Energy (kWh)	
		Voltage	Current	Power	Frequency	Before	After
Fri	2022-04-01	0	0	0	0	0	0
Sat	2022-04-02	0	0	0	0	0	0
Sun	2022-04-03	232.75	9.10	750.56	49.98	899.26	640.37
Mon	2022-04-04	231.16	7.83	803.35	50.17	1215.68	1179.98
Tue	2022-04-05	227.36	6.91	709.84	50.22	1272.20	1272.22
Wed	2022-04-06	227.48	4.60	287.66	50.23	857.12	407.79
Thu	2022-04-07	230.62	8.10	696.73	50.26	1483.23	1473.47
Total Measured for First Week of April 2022						5727.51	4973.83

Note: institutional off-days (e.g., Friday/Saturday) appear as zeros by design.




Figure 7. Weekly total energy before/after optimization for Device 3 (ISO Week 14, 2022).

Table 6. Weekly aggregates for Device 3 in April 2022

Week No.	Average				Weekly Total Energy (kWh)	
	Voltage	Current	Power	Frequency	Before	After
1	230.35	5.22	530.6174	50.14806	5727.51	4973.83
2	232.5078	4.446857	228.5519	50.19547	5716.27	4797.43
3	232.4677	3.695714	420.9319	50.21278	5342.59	4271.72
4	234.1655	5.577857	523.5253	50.24427	9319.08	7064.25
5	236.4911	0	0	50.20386	162.92	0
Total Measured for Month April 2022					26,268.37	21,107.23




Figure 8. Monthly totals in Apr. 2022 for Device 3 (before/after).

Table 7. To date aggregates for Device 3 (Feb 2022–Jan 2023)

Month	Average				Monthly Total Energy (kWh)	
	Voltage	Current	Power	Frequency	Before	After
February	228.65	9.05	1154.69	50.05	19270.58	16592.66
March	229.98	3.95	506.67	49.99	23674.76	18581.10
April	232.82	4.42	393.06	50.20	26268.37	21107.23
May	233.63	1.25	31.92	50.16	3428.18	2196.19
June	230.23	3.15	393.70	50.04	17448.70	13640.70
July	232.01	0.95	71.65	50.06	3443.57	2343.55
September	229.88	3.58	295.39	50.12	20790.76	16368.23
October	232.49	3.29	195.50	50.11	27044.34	16482.86
November	232.16	2.68	340.72	50.10	15463.75	12393.54
December	231.24	2.26	298.13	50.08	17356.28	13544.50
January(2023)	229.24	4.50	784.10	50.09	26194.69	19470.76
Total Measured for whole period (Feb 2022–Jan 2023)					200384	152721.3




Figure 9. Cumulative monthly energy (Feb 2022–Jan 2023) for Device 3.

5.3. Aggregated Total-Consumed-Power reduction ratio

To summarize performance across windows we report a total consumed power reduction ratio defined for any analysis window W by

$$TCPRR(W) = 1 - \frac{\sum_{t \in W} E_{after}(t)}{\sum_{t \in W} E_{before}(t)} \times 100\% \quad (5.1)$$

Where $E(t)$ denotes hourly energy in kilowatt hours. Tables 8 and 9 present the resulting weekly, monthly, and to date ratios for the two devices and Fig. 10 juxtaposes both devices across the three windows. Across all horizons the ratio exceeds 20% and peaking at 25% for Device 1 and about 23.79% for Device 3, which confirms substantial slack in laboratory operation that can be reclaimed through software centric scheduling.

Table 8. The TCP-RR of different periods for d1.

Period	TCP before (kWh)	Hours before	TCP after (kWh)	Hours after	Reduction (%)
Weekly	1888.83	15.62	1563.79	12.31	17.21
Monthly	5558.24	48.66	4158.58	36.00	25.18
To Date	129323.8	491.46	98275.7	355.38	24.01

Table 9. The TCP-RR of different periods for d3.

Period	TCP before (kWh)	Hours before	TCP after (kWh)	Hours after	Reduction (%)
Weekly	5727.51	38.27	4973.83	28.93	13.16
Monthly	26268.37	198.8	21107.23	131.79	19.65
To Date	200384.0	2297.57	154544.51	1348.72	23.79




Figure 10. TCP reduction ratio (weekly, monthly, to-date) for Device 1 and Device 3.

5.4. Statistical validation

Paired two-tailed t-tests and Wilcoxon signed-rank tests applied to synchronized pre and post control samples for phase current, active power, and per sample energy return $p < 10^{-4}$ in every case, which decisively rejects the null hypothesis of random fluctuation [40],[41]. Mean current decreases by roughly thirty percent on both nodes, and effect sizes estimated by Cohen’s d fall in the small range since the controller alters duty cycle rather than instantaneous magnitude. The per sample reduction nevertheless accumulates to multi megawatt hour avoidance over the academic year. The detailed statistics are reported in Tables 10 and 11.

Table 10. Device 1 (Shekhan – SFL): paired-sample statistics.

Variable	μ before	μ after	Δ %	t -stat (p)	Wilcoxon $W(p)$	Cohen’s d
Current [A]	1.72	1.21	-29%	60.6 ($< 10^{-4}$)	1.46×10^{11} ($< 10^{-4}$)	0.12
Power [W]	337	238	-30%	56.2 ($< 10^{-4}$)	1.43×10^{11} ($< 10^{-4}$)	0.11
Energy [kWh][†]	2.81×10^{-3}	1.98×10^{-3}	-30%	56.2 ($< 10^{-4}$)	1.43×10^{11} ($< 10^{-4}$)	0.11

[†] Per-sample energy with a 30 s period: $E = P/1000 \times 30/3600$.

Table 11. Device 3 (Duhok – DFL): paired-sample statistics.

Variable	μ before	μ after	Δ %	t -stat (p)	Wilcoxon $W(p)$	Cohen's d
Current [A]	2.54	1.78	-30%	58.9 ($< 10^{-4}$)	1.40×10^{11} ($< 10^{-4}$)	0.11
Power [W]	512	356	-30%	55.4 ($< 10^{-4}$)	1.38×10^{11} ($< 10^{-4}$)	0.10
Energy [kWh] [†]	4.27×10^{-3}	2.97×10^{-3}	-30%	55.4 ($< 10^{-4}$)	1.38×10^{11} ($< 10^{-4}$)	0.10

[†] Per-sample energy with a 30 s period: $E = P/1000 \times 30/3600$.

5.5. Discussion of results

The scheduler consistently curtails idle demand without degrading productive loading or power quality. Device 1 exhibits larger early stage savings, likely because overnight idling could be suppressed immediately. Device 3 converges over time as boundary tuning improves, which illustrates the adaptive nature of the pipeline through drift triggered retraining and calendar refresh. Because reductions are computed from true power measurements, they translate directly into lower electricity costs and lower associated emissions, both of which are central to campus sustainability objectives. Achievable savings depend on the accuracy of the day-ahead load forecast; larger predictive errors can degrade schedule quality by distorting the expected cost term in (4.1). In practice, the drift-triggered retraining procedure—computed as a rolling 7-day MAPE and fired when it exceeds 10% using a 56-day identification window—mitigated non-stationarity during our study and maintained robust performance.

5.6. Operational implications

At an indicative tariff of one hundred fifty Iraqi dinars per kilowatt hour the combined annual saving of approximately eighty megawatt hours corresponds to about twelve million Iraqi dinars in avoided expenditure. Using a regional grid emission factor of zero point five eight kilograms of carbon dioxide equivalent per kilowatt hour, the same reduction prevents roughly forty six metric tons of carbon dioxide equivalent from entering the atmosphere. Because the scheduler operates entirely in software and executes within a small fraction of the nightly processing window per device per day, these benefits accrue without hardware retrofits, which gives the approach a favorable cost to benefit profile relative to smart plug deployments or panel level upgrades.

6. Conclusions

Universities face persistent pressure to reduce electricity use without compromising core activities. This paper presented an end-to-end intelligent energy-scheduling platform for distributed laboratories at Duhok Polytechnic University that unites high-resolution metering, day-ahead forecasting, and a binary ChOA-based optimizer with calendar feasibility. Over an 11-month horizon the system achieved an average reduction of $\approx 24\%$ —about 80 MWh and 46 t CO₂-e—while remaining operationally light (≈ 180 ms per device per day) and statistically significant (paired tests with $p < 10^{-4}$).

The key contributions are: (1) a deployable, auditable pipeline that closes the loop from 30 s telemetry to executable day-ahead schedules, with explicit logging of model orders and diagnostics; (2) a transparent methodology that discloses SARIMA seasonal orders, a drift trigger (rolling 7-day MAPE $> 10\%$) and a retraining recipe on a 56-day window; (3) a binary ChOA formulation tailored to campus scheduling via a calendar permission mask and a switching regularizer, with penalties normalized by C_{\max} and stated numerically ($\kappa_{\alpha} = 100$,

$\kappa_\beta = 0.05$); and (4) an open, reproducible artifact bundle (cleaned hourly data, serialized models, implementation code, configuration) enabling exact replication and future benchmarking.

Limitations and avenues for future work are clear. Schedules are issued day-ahead rather than in real time, and the evaluation covered two instrumented laboratories rather than the full campus. Achievable savings depend on forecast accuracy; larger day-ahead errors can degrade schedule quality, although the drift-triggered retraining maintained robust performance in our setting. We did not benchmark against alternative binary meta-heuristics in this study; this is a current limitation to be addressed on the same open dataset. Next steps include integrating sub-minute actuation to close the loop in real time, scaling to additional institutes, benchmarking against Binary PSO and Binary Grey Wolf, supporting partial-load control where devices admit intermediate states, exploring hybrid SARIMA–learning forecasters for non-stationary events, and providing a portable calendar-constraint template for rapid adoption.

Taken together, the results show that a software-centric, transparent, and reproducible scheduling approach can deliver persistent, campus-scale energy and emissions savings with minimal operational burden. The released artifacts are intended to catalyze comparative studies and to accelerate evidence-based demand response in higher-education settings within the Kurdistan Region and beyond.

REFERENCES

1. International Energy Agency, Buildings,” 2025. [Online]. Available: <https://www.iea.org/energy-system/buildings> (accessed Sep. 17, 2025).
2. International Energy Agency, *Electricity 2024: Analysis and forecast to 2026*, Jan. 2024. [Online]. Available: <https://www.iea.org/reports/electricity-2024> (accessed Sep. 17, 2025).
3. Kurdistan Regional Government – Ministry of Electricity, “KRG announces new electricity tariffs under Project Runaki,” May 16, 2025. [Online]. Available: <https://gov.krd/ministries/ministry-of-electricity-en/> (accessed Sep. 17, 2025).
4. Ember, “Middle East power sector overview,” Jun. 13, 2025. [Online]. Available: <https://ember-energy.org/insights/research/middle-east-power-sector-overview/> (accessed Sep. 17, 2025).
5. Ember, “Iraq country profile,” Apr. 11, 2025. [Online]. Available: <https://ember-energy.org/countries-and-regions/country-profiles/iraq/> (accessed Sep. 17, 2025).
6. E. C. Quispe, J. F. M. Saldarriaga, and L. A. A. Aponte, “Energy Management Systems in Higher Education Institutions’ Buildings: A Review,” *Energies*, vol. 18, no. 7, p. 1810, 2025. doi: 10.3390/en18071810.
7. M. U. Muqet, S. Soban, and S. U. Khan, “An Overview of Energy Management and Microgrid Control Strategies in University Campuses,” *Energies*, vol. 14, no. 24, p. 8509, 2021. doi: 10.3390/en14248509.
8. S. Lee, J. Park, and S. Kim, “Analysis of Major Temporary Electrical Equipment Energy Consumption in Educational Buildings,” *Sustainability*, vol. 14, no. 17, p. 10783, 2022. doi: 10.3390/su141710783.
9. M. A. Sadeeq, A. O. Musa, and M. A. Shah, “Design and implementation of an energy management system using distributed IoT sensors and Slack,” *Computers & Electrical Engineering*, vol. 111, p. 108775, 2023. doi: 10.1016/j.compeleceng.2023.108775.
10. J. A. Roberson *et al.*, “After-hours Power Status of Office Equipment and Energy Use of Miscellaneous Plug-Load Equipment,” LBNL-53729, 2004. [Online]. Available: <https://eta-publications.lbl.gov/sites/default/files/lbnl-53729.pdf> (accessed Sep. 17, 2025).
11. M. K. Herrlin, “Case Study: JouleX Energy Management (JEM) Solution at Lawrence Berkeley National Laboratory (LBNL),” 2012. [Online]. Available: <https://datacenters.lbl.gov/resources/case-study-jouleX-energy-management-jem> (accessed Sep. 17, 2025).
12. A. Nasir, S. Ahmad, and M. A. Khan, “Impact of COVID-19 lockdown on university campus energy use,” *Energy and Buildings*, vol. 262, p. 111988, 2022. doi: 10.1016/j.enbuild.2022.111988.
13. M. Khishe and M. R. Mosavi, “Chimp Optimization Algorithm,” *Expert Systems with Applications*, vol. 149, p. 113338, 2020. doi: 10.1016/j.eswa.2020.113338.
14. M. Khishe and A. H. Gandomi, “Multi-Objective Chimp Optimizer: An innovative algorithm for multi-objective problems,” *Expert Systems with Applications*, vol. 216, p. 118734, 2023. doi: 10.1016/j.eswa.2022.118734.
15. Y. Wang, H. Gao, and X. Li, “Adaptive Chimp Optimization Algorithm with Chaotic Map,” *The Journal of Supercomputing*, vol. 79, pp. 12589–12617, 2023. doi: 10.1007/s11227-022-04886-6.
16. C. Yue *et al.*, “SCChOA: Hybrid Sine-Cosine Chimp Optimization Algorithm for Feature Selection,” *Computers, Materials & Continua*, vol. 74, no. 1, pp. 2593–2615, 2023. doi: 10.32604/cmc.2023.055059.
17. X. He, Y. Li, and J. Zhang, “Multi-Strategy Enhanced Chimpanzee Optimization Algorithm,” *Applied Sciences*, vol. 15, no. 2, p. 608, 2025. doi: 10.3390/app15020608.
18. M. F. Hamza, B. Modu, and S. Z. Almutairi, “Integration of the Chimp Optimization Algorithm and Rule-Based Energy Management for Microgrid Performance,” *Electronics*, vol. 14, no. 10, p. 2037, 2025. doi: 10.3390/electronics14102037.
19. H. Borousan and M. Hamidian, “Optimal Allocation of Energy Storage Systems in Radial Distribution Networks Using Chimp Optimization Algorithm,” *Sustainability*, vol. 16, no. 19, p. 7904, 2024. doi: 10.3390/su16197904.

20. S. Al-Dahidi *et al.*, “Enhancing solar photovoltaic energy production prediction using ML models tuned with the Chimp Optimization Algorithm,” *Scientific Reports*, vol. 14, p. 18583, 2024. doi: 10.1038/s41598-024-68595-8.
21. Z. Zhang, H. Liu, and P. Wang, “A Dual-Adaptive Stochastic Reinforcement Chimp Optimization Algorithm,” *Scientific Reports*, vol. 14, p. 6910, 2024. doi: 10.1038/s41598-024-52820-5.
22. Z. Zhang, C. Xu, and P. Wang, “Evolving Chimp Optimization Algorithm Using Quantum Computing,” *Journal of Computational Design and Engineering*, vol. 11, no. 5, pp. 143–158, 2024. doi: 10.1093/jcde/qwae010.
23. E. Y. Alhawsawi, I. Alsaïdan, and A. Al-Sumaiti, “A comprehensive review of existing and pending university campus microgrids,” *Energies*, vol. 17, no. 10, p. 2425, 2024. doi: 10.3390/en17102425.
24. S. Alsamraee and S. Khanna, “High-resolution energy consumption forecasting of a university campus power plant based on advanced machine learning techniques,” *Energy Strategy Reviews*, vol. 60, p. 101769, 2025. doi: 10.1016/j.esr.2025.101769.
25. D. Ashtar, A. Homairfar, and M. A. Hannan, “Hybrid forecasting for sustainable electricity demand using SARIMAX–LSTM and sequence models,” *Sustainability*, vol. 17, no. 16, p. 7192, 2025. doi: 10.3390/su17167192.
26. H. Jiang, S. Zhang, and Y. Wang, “Online SARIMA applied for short-term electricity load forecasting,” *Applied Intelligence*, vol. 53, pp. 13257–13273, 2023. doi: 10.1007/s10489-022-04265-6.
27. S. Tijjani, O. T. Akinsanmi, and A. H. Ganiyu, “An enhanced particle swarm optimization with position update for optimal feature selection,” *Expert Systems with Applications*, vol. 236, p. 123337, 2024. doi: 10.1016/j.eswa.2023.123337.
28. J. Y. Khaseeb, M. Islam, and R. A. R. Ghazali, “Improved binary grey wolf optimization approaches for feature selection,” *Applied Sciences*, vol. 15, no. 2, p. 489, 2025. doi: 10.3390/app15020489.
29. R. Korab, P. Čič, and L. Fojtů, “Optimal Scheduling of Energy Storage and Shiftable Loads Considering Demand Response in Residential Buildings,” *Energies*, vol. 17, no. 21, p. 5264, 2024. doi: 10.3390/en17215264.
30. T. Domínguez-Bolaño *et al.*, “An IoT system for a smart campus: Challenges and solutions,” *Sustainable Computing: Informatics and Systems*, vol. 45, p. 100880, 2024. doi: 10.1016/j.suscom.2024.100880.
31. D. Hercog and B. Gergič, “Design and Implementation of ESP32-Based IoT Devices,” *Sensors*, vol. 23, no. 13, p. 5997, 2023. doi: 10.3390/s23135997.
32. A. Manowska, A. Bocewicz, and Z. Banaszak, “The Use of the MQTT Protocol in Measurement, Monitoring and Control Systems,” *Electronics*, vol. 12, no. 1, p. 17, 2022. doi: 10.3390/electronics12010017.
33. A. Franco *et al.*, “Energy Savings in University Buildings: The Potential Role of IoT Sensor Networks,” *Sustainability*, vol. 17, no. 1, p. 111, 2024. doi: 10.3390/su17010111.
34. A. Limane and F. Gagnon, “Why Chaos Works in Metaheuristics: A Survey and New Insights,” *Information Sciences*, vol. 665, pp. 119140, 2025. doi: 10.1016/j.ins.2024.119140.
35. Y. Yufeng, Z. Yong, and C. Yi, “Intelligent day-ahead optimization scheduling for multi-community energy supply with CCHP,” *Frontiers in Energy Research*, vol. 11, Art. no. 1349194, 2024. doi: 10.3389/fenrg.2023.1349194.
36. International Energy Agency, *Energy and AI*, 2025. [Online]. Available: <https://www.iea.org/reports/energy-and-ai> (accessed Sep. 17, 2025).
37. U.S. DOE Better Buildings, “A Primer on Organizational Use of Energy Management and Information Systems (EMIS),” 2016. [Online]. Available: <https://betterbuildingssolutioncenter.energy.gov/> (accessed Sep. 17, 2025).
38. A. Shehabi *et al.*, “United States Data Center Energy Usage Report,” LBNL, 2016. [Online]. Available: https://www.iea-4e.org/wp-content/uploads/publications/2016/06/05j_-_LBNL_-_US_Data_Centres_Energy_USE.pdf (accessed Sep. 17, 2025).
39. S. Gheouany, H. Ouadi, F. Giri, and S. El Bakali, “Optimal Day-Ahead Active and Reactive Power Management in Residential Buildings using Particle Swarm Optimization,” *IFAC-PapersOnLine*, vol. 58, no. 13, pp. 170–175, 2024. doi: 10.1016/j.ifacol.2024.07.478.
40. W. S. Gosset (“Student”), “The Probable Error of a Mean,” *Biometrika*, vol. 6, pp. 1–25, 1908.
41. H. B. Mann and D. R. Whitney, “On a Test of Whether One of Two Random Variables Is Stochastically Larger than the Other,” *Annals of Mathematical Statistics*, vol. 18, pp. 50–60, 1947.