2026 Design, Automation & Test in Europe Conference (DATE 2026)

Artemis: Co-Simulation of Power Microgrids and Energy-Aware Cloud Data Centers

Mattia Tibaldi
Politecnico di Milano
mattia.tibaldi@polimi.it

Sara Vinco
Politecnico di Torino
sara.vinco@polito.it

Christian Pilato
Politecnico di Milano
christian.pilato@polimi.it

***Abstract*—The growing demand for power to support new cloud services raises the question of how to power future data center infrastructures. A power microgrid and cloud simulator that can act as a unified digital twin of these new infrastructures is crucial for studying emerging scenarios. In this article, we propose Artemis, a co-simulation environment for power microgrids and cloud data centers. Artemis extends the combination of the CloudSim Plus simulator and the Amethyst virtual machine allocation and migration policy with a generalized power microgrid model. Ultimately, Artemis enables the study of modular power microgrids with custom electrical policies and returns performance metrics and visualizations of the data center's status under observation.**

I. Introduction

With the emergence of new technologies based on artificial intelligence (AI) and the growing attention in sustainable development, the demand for sustainable data centers and sustainable energy sources is rapidly increasing [1], [2]. To avoid power outages and reduce electricity costs and carbon footprint, many data centers are supported by power *microgrids*, which provide renewable power generation and power storage, thereby enhancing energy efficiency, sustainability, and reliability [3] (see Figure 1). Examples of power sources adopted by such microgrids include traditional solar farms, wind turbines, and thermal plants, as well as more recent solutions such as on-site small modular reactors (SMRs) [4] and sand batteries [5]. The combination of these electrical technologies makes modern power microgrids complex to study [6]: determining which power resources to use and which electrical policies to apply is challenging, especially when connected to a data center with variable workloads.

The literature is rich in simulators that consider power microgrids and data centers as distinct entities [7]–[9]. However, few tools combine these two aspects in a unified simulation system [10]–[13], and, to the best of our knowledge, none of them allows us to study a modular power microgrid with custom electrical policies and observe how the data center reacts and vice versa. Having complete control over the power microgrid is essential for studying how different configurations affect the data center and for answering questions such as: *What happens if we add a wind turbine to our system? What happens if we remove the storage and use a policy that constantly draws on fossil fuels?*

In this paper, we present Artemis, a power microgrid and cloud co-simulation environment developed to help answer these questions. Artemis extends CloudSim Plus [14] and the Amethyst virtual machine (VM) migration and allocation algorithm [15] for data center VM orchestration with a generalized power microgrid model. The combination of these components enables a unified, extensible simulation environment in which the data center responds to the power microgrid's electrical policy and the runtime workload. The contributions of this article are:

- A framework that unifies power microgrid and data center models in a single simulated environment;
- An evaluation of the proposed framework with different electrical policies, microgrid configurations, and data center workloads.

The Artemis framework is available as open source [16]. The rest of the paper is organized as follows. Section II outlines the proposed framework, which is detailed in Section III. Section IV presents the simulation results. Finally, Section V summarizes the state of the art and Section VI draws our concluding remarks.

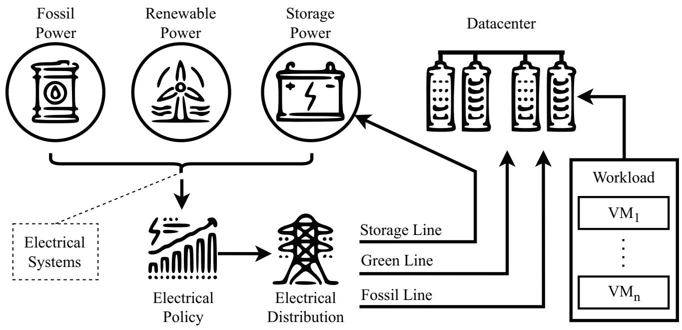

Fig. 1. Heterogeneous electrical systems compose the power microgrid. The electrical policy determines the power amount drawn from each resource and injects it onto the electrical distribution power lines leading to the data center.

II. Artemis Overview

Artemis is a framework built on CloudSim Plus [14] and Amethyst [15]. Given a power microgrid configuration, an electrical management policy, a data center configuration, and a workload, it simulates the infrastructure and evaluates its overall performance.

The power microgrid is a localized energy system that can generate, store, and distribute electricity, either independently or in coordination with the principal electrical grid. It is composed of several *electrical systems* (typically photovoltaic panels and wind turbines) and storage units (i.e., *batteries*)

978-3-9826741-1-7/DATE26/© 2026 EDAA

managed by an *electrical policy* that allocates available power to satisfy a given power demand [17], [18].

We design the power microgrid to serve a data center, configured as several *hosts* that can change their configuration at runtime to adapt to the data center's and power microgrid's needs, e.g., due to multi-variant hardware. We consider two possible *computing modes*: low-power (e.g., use the hardware variant that consumes the least power) and max-quality (e.g., use the hardware variant that guarantees the highest quality-of-service (QoS)). Host workload is modeled as VMs and Cloudlets (e.g., tasks executed on the VMs) [14].

Figure 2 summaries the architecture of Artemis. The baseline is implemented in CloudSim Plus (blue), on top of which Amethyst builds the data center orchestration model (purple) [15]. Artemis extends such a simulation with the power microgrid model (yellow), so that power configuration and management policy influence the data center:

- Amethyst computes power overall request at time t, based on the data center's upcoming activity (e.g., the next scheduling instance t in the simulation);
- the microgrid estimates the available power from both renewable and non-renewable sources, and applies an electrical policy that can prioritize renewable sources, balance storage usage, or apply any desired strategy for the power side of the system;
- the output of the policy is the allocation of power demand onto the different classes of power providers;
- based on the power solution, the Amethyst policy orchestrates the data center, determining active hosts (e.g., active servers), computing modes, and VM allocation.

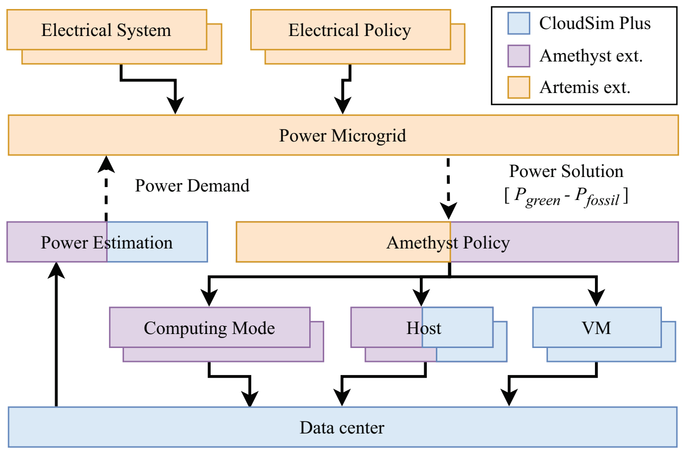

Fig. 2. Proposed framework. The figure highlights the Artemis extension (yellow) and the relations with CloudSim Plus [14] and Amethyst [15]. The two-colored blocks indicate that the blocks, originally present in CloudSim Plus and Amethyst, have been extended in the current work.

III. Artemis Simulation Environment

This section details the Artemis simulation environment. Amethyst manages the simulation, exploiting the CloudSim event-driven scheduling to evaluate VM allocation and migration at fixed time steps. We extend Amethyst to concurrently invoke the power microgrid simulation, sending the calculated data center power demand. After that, the microgrid returns a power solution according to the active electrical policy.

A. Data Centers Orchestration Model

Amethyst is a VM allocation and migration algorithm that aims to reduce carbon emissions in data centers through VM intelligent management. It is an extension to CloudSim Plus for power-aware data center orchestration.

The algorithm takes available power as input, distinguishing between renewable and fossil fuel sources, and computes two dynamic consumption thresholds (e.g., lower-bound and upper-bound thresholds). Its goal is to keep each host's consumption within the two thresholds.

Amethyst will attempt to configure the hosts in low-power mode and consolidate the VMs before pausing the execution of some workloads or, in extreme cases, shutting down the machine if the host consumption exceeds the power thresholds. Conversely, it will configure the hosts to max-quality mode if power is sufficient. This configuration takes into account the available power input by distinguishing green and fossil fuels (to consider carbon emissions in the policy). However, power inputs are fixed temporal traces; thus, the impact of power aspects on the data center is limited.

1) Data Center Power Estimation: In this work, the main consumers are the hosts in the data center. Amethyst estimates the data center's power demand, $PD_{dc}(t)$ as:

$$
\begin{aligned}
PD_h(t) &= SP_h(t) + U_h(t) \cdot (MP_h(t) - SP_h(t)) \\
PD_{dc}(t) &= \sum_{h}^{Host} A_h(t) \cdot PD_h(t)
\end{aligned} \tag{1}
$$

where $U_h(t)$ is host's resources percentage used, $PD_h(t)$ is the power demand of the host h at time t, $SP_h(t)$ and $MP_h(t)$ are the static and maximum power of host h at time t, respectively, and $A_h(t)$ indicates whether host h is active at time t. SP_h and MP_h may vary over time depending on the computing mode.

B. Power Microgrid Model

The goal of Artemis is to support Amethyst with a simulatable implementation of the available power. The power generation becomes a design parameter to support data center simulation, with the possibility to compare alternative microgrid configurations. Several classes of electrical systems, each with a distinct role in generating, storing, converting, distributing, and controlling power, form the microgrid [19]. Power sources actively generate power, e.g., from solar irradiance or wind, while power storage units (e.g., batteries) balance supply and demand and provide power backup.

Given that the goal of a microgrid for server farms is to serve continuous 24/7 demand with very high reliability and redundancy, the microgrid will include several types of power sources and storage devices. We describe each electrical system by its produced power, $P_{es}(t)$.

1) Solar Panel Model: We model the solar panel as presented by Dubey et al. [20]. The output of a solar panel varies as a function of the global horizontal irradiance $GHI(t)$ and its area a and temperature $K(t)$:

$$\begin{aligned} EF_s(t) &= cp \cdot [1 - ct \cdot (K(t) - nt)] \\ P_{es}(t) &= EF_s(t) \cdot GHI(t) \cdot a \end{aligned} \tag{2}$$

where $EF_s(t)$ is panel efficiency calculated from the constants: panel cell efficiency factor cp, panel loss percentage per degree ct, and panel nominal temperature nt.

2) Wind Turbine Model: We model the wind turbine as proposed by Kim et al. [21]. The output of a wind turbine varies as a function of air pressure $AP(t)$, air temperature $K(t)$, wind speed $V(t)$, and rotor blade radius br:

$$\begin{aligned} AD(t) &= \frac{AP(t)}{\rho \cdot K(t)} \\ P_{es}(t) &= \frac{1}{2} \cdot cp \cdot \pi \cdot br^2 \cdot AD(t) \cdot V(t)^3 \end{aligned} \tag{3}$$

where ρ is the gas constant for dry air, $AD(t)$ represents air density, and cp is a constant power coefficient of the turbine.

3) Thermal Plant Model: Simulating a thermal power plant is complex as it is composed of multiple sub-systems (e.g., boiler, turbine, and steam transmission), each with dedicated models [22]. Modeling such a solution in detail is beyond the scope of this article. Artemis models the thermal power plant in a simplified way [23]. Power plant output is given by a constant value tp subject to a power plant efficiency $EF_t(t)$ as:

$$P_{es}(t) = tp \cdot EF_t(t) \tag{4}$$

4) Storage Model: Artemis models the storage battery as a simplified discrete-time Markov chain with $N + 1$ states [24]. The state number N corresponds to the number of charge units available in the battery. The charging or discharging of a charge unit determines the movement along the chain: a charge unit is consumed with probability q or recovered with probability $1 - q$, and the battery is empty when the system reaches the absorbing state 0. In Artemis, we replace the probability q of taking an action with a simulated action determined by the electrical policy. The storage is a special electrical system that can also act as a power consumer when discharged. Therefore, in addition to P_{es}, we also model its power demand PD_{es}. Note that the power microgrid does not necessarily have to satisfy this power demand, but it can defer it over time. The resulting model is as follows:

$$\begin{aligned} P_{es}(t) &= \begin{cases} pu \cdot EF_b(t) & cu < N + 1 \\ 0 & otherwise \end{cases} \\ PD_{es}(t) &= cu \cdot pu \cdot EF_b(t) \end{aligned} \tag{5}$$

where pu is the power contained in a charge unit, $EF_b(t)$ is the efficiency of the battery at time t and cu is the number of unit charges consumed.

C. Electrical Policy Model

The electrical policy is a model that takes as input the set of electrical components in the power microgrid $ES(t)$ and the data center power demand $PD_{dc}(t)$ and returns an power solution $S(t) = \{s^{c0}(t), \ldots, s^{cn}(t)\}$ representing the set of power quantities per considered resource category. A resource category $c \in C$ is an abstraction for defining complex decisions on the available power, e.g., the type of power produced, such as renewable or fossil fuel.

Each electrical policy starts by calculating the power available at time t by resource category:

$$P_a^c(t) = \sum_{es}^{ES(t)} P_{es}^c(t) \quad \forall c \in C \tag{6}$$

where $P_{es}^c(t)$ is the power produced by electrical system es at time t in category c. The power solution $S(t)$ is the result of the integer linear programming (ILP) model at the time t, subject to the custom set of constraints $R(t)$:

$$\begin{aligned} max \sum_c^C s^c \quad & subject\ to \\ \sum_c^C s^c &\le PD_{dc} \\ s^c &\le P_a^c \quad \forall c \in C \\ s^c\ subject\ to\ r & \quad \forall r \in R \end{aligned} \tag{7}$$

i.e., the ILP aims at maximizing the overall available power ($max \sum_c^C s^c$), imposing that the power selected per each category is at most the power produced by such power source ($s^c \le P_a^c$), and that the selected amount of power per each power category is determined by the constraints $R(t)$. The way we define the set of such constraints defines the electrical policy (e.g., always prefer renewable sources, always ensure a certain level of state of charge in the battery, and so on). The output $S(t)$ of the ILP defines the amount of power used to satisfy the data center demand at time t.

As additional output, the power solution $S(t)$ includes a sum of the not assigned power (i.e., $\{s^c \in S : s^c < P_a^c\} \neq \emptyset$). Amethyst uses this surplus to activate deferred tasks, including battery charging (e.g., satisfy the storage $PD_{es}(t)$).

D. Electrical Distribution Model

Artemis implements a basic model for the electrical distribution system, consisting of the set of power destinations D, which represent separate physical lines carrying power to end users. Amethyst assumes in its model that there are two distinct power lines, green and fossil, arriving at the data center [15]. Artemis represents an electrical distribution as a function $f_{ed} : C \to D$, where C is the set of implemented categories described in Section III-C. As a result, the components s^c of a generic power solution S are grouped into the correct destination. This modeling is essential to enable proper integration between Amethyst and the power microgrid, allowing us to experiment with elaborate policies while maintaining output consistent with what Amethyst expects.

IV. Artemis Simulation Results

We developed Artemis in Java as an extension of CloudSim Plus and Amethyst. To demonstrate the framework's capabilities, we evaluated four electrical policies in different case

TABLE I
POWER GRID CONFIGURATION. N IS THE NUMBER OF INSTANCES.

Grid	**Solar P.**		**Wind T.**		**Storage**		**Thermal P.**	
	a m^2	N -	br m	N -	$cu \times pu$ Wh	N -	tp kW	N -
G.1	2.2	20	1.5	5	2,000	2	4	1
G.2	2.2	20	0.0	0	2,000	2	4	1
G.3	2.2	20	1.5	5	2,000	6	4	1
G.4	2.2	100	1.5	50	2,000	60	60	1

studies. We defined each policy as a mapping from the set $C = \{green, fossil\}$ of categories to the set $R(t)$ of ILP constraints.

Definition 1: Let **green-first (GF)** be the policy that provides the maximum power capacity, prioritizing power generated from renewable resources; $R(t) = \{s^{fossil} \leq PD_{dc} - P_a^{green}\}$.

Definition 2: Let **fixed-fossil (FF)** be the policy that provides a maximum amount of power pf from the fossil power line and manages the rest with renewable resources and battery storage; $R(t) = \{s^{fossil} \leq pf,\ s^{green} \leq PD_{dc} - pf\}$.

Definition 3: Let **cost-aware (CA)** be the policy that provides a mix of power based on the hourly cost of fossil energy. We develop this policy by considering electricity pricing schemes that increase tariffs during peak-demand hours (e.g., from 4 pm to 7 pm [25]). Specifically, from 12 am to 5 am, it uses power generated by wind turbines to recharge the battery and manage the data center with fossil fuels; from 6 am to 4 pm, it uses a 50/50 mix of fossil and green power; from 4 pm to 7 pm, it uses renewable energy and batteries to reduce emissions, and fossil fuels in the remaining time.

$$R(t) = \begin{cases} \{s^{green} = 0\} & 12\ a.m < t < 5\ a.m \\ \{s^{green} \leq 0.5 \cdot PD_{dc}, & \\ \ s^{fossil} \leq 0.5 \cdot PD_{dc}\} & 6\ a.m < t < 4\ p.m \\ \{s^{fossil} = 0\} & 4\ p.m < t < 7\ p.m \\ \{s^{green} = 0\} & otherwise \end{cases}$$

Definition 4: Let **grid-relief (GR)** be the policy that provides a maximum amount of power pg from renewable resources and manages peaks with only battery storage; $R(t) = \{s^{green} \leq pg,\ s^{fossil} = 0\}$.

A. Electrical Policy Comparison in Different Seasons

The goal is to demonstrate how Artemis can compare different electrical policies across various seasons. This type of study is useful for evaluating how to develop effective policies that maintain a constant service rate.

1) Simulation Setup: We configured the data center as described by Ye et al. [12]. The data center consists of 10 hosts, each with a computational power of 2,000 million instructions per second (MIPS) and a power consumption (Watts) of $\{SP_h = 218.0,\ MP_h = 350.0\}$ in max-quality and $\{SP_h = 196.2,\ MP_h = 315.0\}$ in low-power. The workload consists of 20 VMs with a stochastic $U_h(t)$ utilization model. The power microgrid configuration is *G.1*, as shown in Table I.

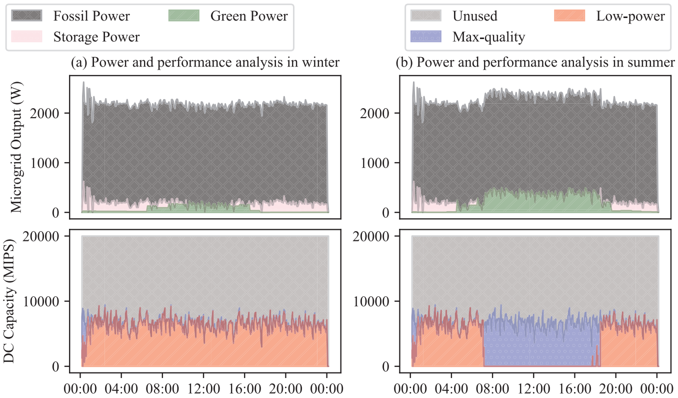

Fig. 3. Data center and microgrid behavior when applying FF policy on a winter and summer day.

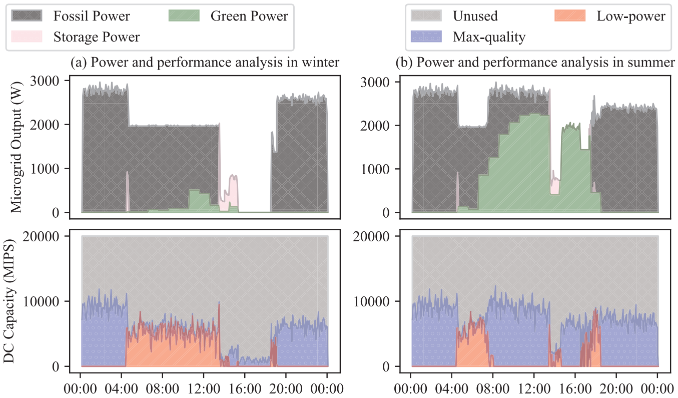

Fig. 4. Data center and microgrid behavior when applying CA policy on a winter and summer day.

Since Ye et al. model only battery usage, the storage values are faithful to the original. Finally, we consider the weather trace for two days in 2018 [26], one winter and one summer.

2) Simulation Results: Figures 3 and 4 present the simulation results for the FF and CA policies. We limit the discussion to these two policies due to space constraints, while a full design space exploration is provided in Section IV-B. Our analysis focuses on the consistency of the data center's power supply and the resulting QoS (i.e., the percentage of time it operates at maximum quality).

In Figure 3 with a $pf = 2$ kW, the data center consumes more than the set maximum value and attempts to manage peaks using battery and renewables. However, on a winter day, FF is ineffective, and the data center continuously operates in low-power mode, resulting in suboptimal QoS. In summer, conditions improve only during daylight, when the panels enable the data center to operate at peak performance. Figure 4 shows the outcomes under the same conditions when applying the CA policy. On a winter day, renewable power production is insufficient in the afternoon hours. So, Amethyst completely shuts down the data center. On a summer day, QoS improves, but due to a cloudy period, the batteries kick in. In this case, Amethyst pauses some workload executions, as available resources cannot cover the entire load.

In conclusion, the CA policy poses significant service issues

TABLE II

RESULTS COMPARISON OVER ONE MONTH OF EXECUTION. HIGHLIGHTED THE PREFERABLE AND NON-PREFERABLE STATISTICS.

Grid	Policy	Demand kWh	Green %	Fossil %	Storage %	Charge %	Price €	Emission $KgCO_2e$	SLAV $x10^{-7}$	Mig. -	Change -	QoS %	Workload %
G.1	GF	2,061	35.43	64.57	0.00	100.00	340.30	466.02	0.0	0	0	100.00	100.00
	FF	1,074	11.80	87.39	0.81	90.71	248.07	323.18	1,339.3	19	12,700	38.99	99.99
	GR	462	86.53	0.00	13.47	52.28	0.00	5.62	13.9	9	1,403	76.52	39.66
	CA	1,244	28.54	69.89	1.57	94.47	209.62	302.92	24.2	6	3,426	74.34	42.48
G.2	GF	1,877	32.17	67.83	0.00	100.00	319.89	444.82	0.0	0	0	100.00	100.00
	FF	1,370	3.88	96.11	0.01	99.13	338.25	451.69	746.1	20	12,486	46.41	99.99
	GR	368	98.30	0.00	1.70	52.84	0.00	5.08	379.4	7	1,161	86.49	41.17
	CA	1,324	25.29	72.12	2.59	92.16	229.37	331.79	9.4	5	2,598	58.58	36.45
G.3	GF	1,653	39.77	60.23	0.00	100.00	256.66	350.39	0.0	0	0	100.00	100.00
	FF	1,476	3.98	95.93	0.09	99.99	360.82	485.78	640.3	18	5,316	45.49	100.00
	GR	410	77.56	0.00	22.44	56.12	0.00	4.47	72.7	4	4,394	64.28	41.92
	CA	1,240	25.60	74.22	0.18	97.79	218.48	319.75	5.7	11	9,852	67.83	96.56

with the G.1 microgrid configuration. The choice falls on FF even if the final QoS is low.

B. Power Microgrid Design Exploration

Artemis can evaluate different power microgrid configurations and electrical policies over the long term. This study is essential for building an optimal microgrid power solution.

1) Simulation Setup: The setup is the same as in Section IV-A, but with different power microgrid configurations (see Table I) and a more extended execution period of one month. Weather traces have been collected in spring, to allow for greater atmospheric variability [26]. To calculate the energy price, we use the data tabulated in [25], while to estimate emissions, we use the data reported in [15].

2) Simulation Results: Table II shows the metrics obtained from the simulation, in terms of total energy consumed (broken down into green, fossil, and storage rates), average battery charge percentage, fossil energy cost, overall CO_2 emissions, service level agreement violations (SLAVs) (calculated as in [15]), number of migrations, number of mode changes, QoS, and percentage of completed workload.

The goal is to find a solution that reduces emissions while still enabling workload computation with good QoS and low SLAVs. Starting from configuration G.1, we conclude that the most promising solutions are GR and CA, as they achieve low emissions, low SLAV, and QoS above 70%. Unfortunately, the computed workload percentage is low, indicating that the power microgrid is undersized.

The subsequent experiment thus consists of modifying the power microgrid configuration to evaluate the impact that individual electrical components have on the data center. The G.2 configuration eliminates the wind turbines. The result is an immediate performance degradation. The data center switches to fossil fuels, and the computed workload drops further, indicating that the problem lies in managing early-morning time slots, when the turbines operate more frequently, and the data center struggles to maintain constant activity.

In G.3, we reintegrate the turbines and increase the number of batteries, enabling the data center to manage a power service drop. Finally, the G.3 CA configuration emerges as a mediator between the metrics under consideration.

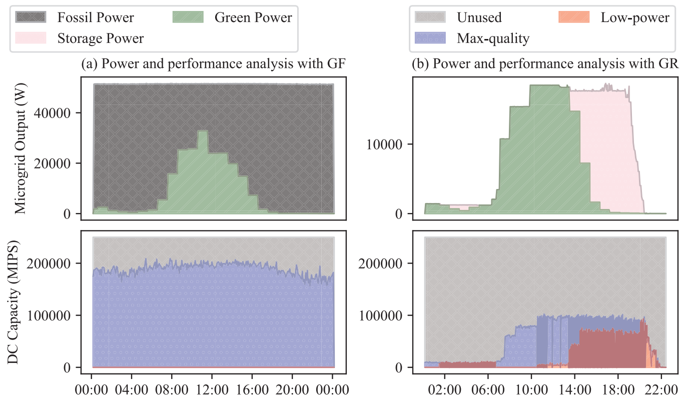

Fig. 5. PlanetLab and microgrid behavior when applying GF and GR policies.

C. Artemis Scalability

We analyze the PlanetLab data center [27] to demonstrate how Artemis manages extensive simulations.

1) Simulation Setup: PlanetLab consists of 800 heterogeneous hosts evenly split between HP G4 and HP G5 models [27]. The HP G4 host has 1,860 MIPS of computing power and power consumption of $\{SP_h = 74.4,\ MP_h = 105.3\}$ in low-power and $\{SP_h = 86.0,\ MP_h = 117.0\}$ in max-quality mode. The HP G5, on the other hand, has 2,660 MIPS and a power consumption of $\{SP_h = 86.4,\ MP_h = 121.5\}$ in low-power and $\{SP_h = 94.0,\ MP_h = 135.0\}$ in max-quality mode. The workload consists of 1,052 VMs as described in [27]. Actual daily traces define the resources' utilization $U_h(t)$. We adopt the power microgrid $G.4$ (see Table I).

2) Simulation Results: Figure 5 shows the results with GF and a GR policy. In the first case, the data center computes the entire workload. In the second case, computation is highly discontinuous, and only 53.77% of the workload is completed. These two cases allow us to discuss the simulation execution times, which depend on the simulated events. In particular, VM allocation events have more weight because they require cycling through hosts, and the power microgrid must generate a solution at each iteration. We estimate these events as:

$$N_{ae} \sim N_{vm} \cdot (1 + P_f \cdot \frac{S_{tl}}{S_c}) + N_m \tag{8}$$

TABLE III
STATE-OF-THE-ART SUMMARY. WE COMPARE THE MAIN SIMULATOR FEATURES CONSIDERED IN THE ARTICLE TO OUR PROPOSED SOLUTION.

Works	Power Grid Simulation			Cloud Simulation				
	Elect. System	*Elect. Policy*	*Elect. Distribution*	*Hardware*	*Network*	*Power*	*Sch. Policy*	*All. & Mig. Policy*
[7]	-	-	-	advanced	basic	basic	basic	basic
[14]	-	-	-	basic	medium	basic	medium	advanced
[8]	-	-	-	medium	advanced	medium	basic	advanced
[9]	-	-	-	medium	medium	advanced	advanced	advanced
[28]	-	advanced	advanced	-	-	-	-	-
[29]	advanced	advanced	medium	-	-	-	-	-
[10]	medium	basic	-	basic	-	medium	-	-
[11]	-	medium	-	-	-	basic	-	-
[12]	basic	-	basic	basic	medium	basic	medium	advanced
[13]	medium	-	basic	basic	-	basic	-	-
Artemis	advanced	advanced	basic	medium	medium	basic	medium	advanced

Where N_{vm} is the number of VMs, P_f is the failure rate computed as one minus the computed workload percentage, S_{tl} is the simulation time limit (e.g., one day), S_c is the simulation clock, and N_m is the number of migrations. In the simulation, we set a clock of 300 seconds, and in both cases, there are no migrations. For GF, this involves 1,052 events, whereas for GR it is 140,968 ($\sim$4.5 hours on an Intel i5-11600 K).

In conclusion, Artemis handles extensive simulations with a running time that increases slowly with the number of VMs but rapidly with the failure rate and the simulation time limit.

V. RELATED WORK

We categorize previous works based on their functionality and if they consider the power microgrid and cloud in a unified simulation. Table III compares the features with our solution.

A. Power-Aware Cloud Simulation

Simulating cloud infrastructures with a focus on power consumption is crucial for designing energy-efficient data centers and evaluating resource allocation strategies. Several simulators exist to address this need. HET-SIM [7] targets heterogeneous hardware by supporting accelerators such as FPGA and GPU, but its power and scheduling models are limited, and it lacks dynamic workload allocation. CloudSim Plus [14], in contrast, is a flexible event-driven framework for simulating IaaS clouds. It is well suited to performance and resource-usage evaluation, though it places less emphasis on detailed power modeling. GreenCloud [8] focuses directly on energy-aware strategies, provides accurate component model consumption estimates, and leverages virtualization to reduce overall power consumption. Finally, OpenDC 2.0 [9] includes facility components (e.g., cooling systems) and integrates real-world orchestrators such as Kubernetes and OpenStack into the simulation. All such solutions, however, do not support a dynamic model of the power infrastructure and thus have limited capacity to incorporate power optimization into their algorithms.

B. Power Microgrid Simulation

Modeling and simulating power microgrids is a long-established research area that has been extensively studied from multiple perspectives, including operational costs, geographic distribution, and control [30]–[33]. Such works, however, consider power consumers only as inputs for the power management policy. ESPSim [28] is a large-scale power flow simulator, optimized to solve power flow problems. But it does not incorporate detailed physical models of electrical components. In contrast, Rendel et al. [29] focused on estimating market costs but lacked flexibility in configuration and omitted analysis of alternative distribution or operational policies.

C. Co-simulation of Power Microgrid and Cloud

Cloud and power microgrid simulators typically address separate concerns, yet data centers require integrated modeling of computing workloads and energy supply. Co-simulation approaches aim to bridge this gap [34]. ReRack [10] models renewable-powered data centers by combining an annual cost simulator with a genetic optimizer for source selection, though workloads are represented only by traces. Wang et al. [11] adopt a two-stage model where microgrid operators balance loads and data centers minimize costs via an ILP formulation. This solution provides economic insights but lacks detailed system modeling. Power Supply CloudSim [12] extends CloudSim [35] with power supply features, particularly UPS, to study demand–response strategies. Unfortunately, they support only storage batteries. Zhang et al. [13] propose a time-series simulation of wind–battery–grid integration for data centers, optimizing costs across multiple sources. The microgrid model remains fixed and non-extendable. Overall, current co-simulation tools highlight the benefits of integrating cloud and microgrid models but remain limited by simplified workloads, rigid microgrid configurations, and limited policy flexibility.

VI. CONCLUSION AND FUTURE WORK

We presented Artemis, a co-simulation framework for studying the interaction between a data center and a power microgrid. We evaluate Artemis on several case studies. Artemis enables in-depth exploration of microgrid design and data center monitoring, generating a comprehensive view of the entire system that allows for the analysis of more efficient solutions.

We plan to exploit Artemis to implement the digital twin of a new physical infrastructure, implementing Amethyst as an autonomous orchestration agent.

References

[1] P. Devarakota, N. Tsesmetzis, F. O. Alpak, A. Gala, and D. Hohl, "AI and the net-zero journey: Energy demand, emissions, and the potential for transition," 2025. [Online]. Available: https://arxiv.org/abs/2507.10750

[2] The International Energy Agency (IEA). (2025) Energy demand from AI. [Online]. Available: https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai

[3] L. Yu, T. Jiang, and Y. Zou, "Distributed real-time energy management in data center microgrids," *IEEE Transactions on Smart Grid*, vol. 9, no. 4, pp. 3748–3762, 2018.

[4] E. Waltz. (2024) Big tech backs small nuclear Google and Amazon invest in small modular reactors to power data centers. [Online]. Available: https://spectrum.ieee.org/nuclear-powered-data-center

[5] Polar Night Energy. (2024) What is a sand battery? [Online]. Available: https://polarnightenergy.com/news/what-is-a-sand-battery/

[6] A. M. Amani and M. Jalili, "Power grids as complex networks: Resilience and reliability analysis," *IEEE Access*, vol. 9, pp. 119 010–119 031, 2021.

[7] E. Kanellou, N. Chrysos, and A. Bilas, "A flexible datacenter simulator," *Procedia Computer Science*, vol. 136, pp. 72–81, 2018, 7th International Young Scientists Conference on Computational Science, YSC2018, 02-06 July2018, Heraklion, Greece. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1877050918315448

[8] L. Liu, H. Wang, X. Liu, X. Jin, W. B. He, Q. B. Wang, and Y. Chen, "Greencloud: a new architecture for green data center," in *Proceedings of the 6th International Conference Industry Session on Autonomic Computing and Communications Industry Session*, ser. ICAC-INDST '09. New York, NY, USA: Association for Computing Machinery, 2009, p. 29–38. [Online]. Available: https://doi.org/10.1145/1555312.1555319

[9] F. Mastenbroek, G. Andreadis, S. Jounaid, W. Lai, J. Burley, J. Bosch, E. van Eyk, L. Versluis, V. van Beek, and A. Iosup, "Opendc 2.0: Convenient modeling and simulation of emerging technologies in cloud datacenters," in *2021 IEEE/ACM 21st International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*, 2021, pp. 455–464.

[10] M. Brown and J. Renau, "Rerack: power simulation for data centers with renewable energy generation," *SIGMETRICS Perform. Eval. Rev.*, vol. 39, no. 3, p. 77–81, Dec. 2011. [Online]. Available: https://doi.org/10.1145/2160803.2160865

[11] H. Wang, J. Huang, X. Lin, and H. Mohsenian-Rad, "Exploring smart grid and data center interactions for electric power load balancing," *SIGMETRICS Perform. Eval. Rev.*, vol. 41, no. 3, p. 89–94, Jan. 2014. [Online]. Available: https://doi.org/10.1145/2567529.2567556

[12] G. Ye, H. Wang, Y. Li, Z. Wang, and F. Gao, "Power Supply Cloudsim: An extension of CloudSim for modeling and simulation of power supply devices in cloud data centers," in *2022 IEEE/IAS Industrial and Commercial Power System Asia (I&CPS Asia)*, 2022, pp. 1620–1626.

[13] D. Zhang, K. Liang, and P. You, "Study on data center wind-storage-load-grid system optimization considering time-of-use prices," in *2024 6th International Conference on Power and Energy Technology (ICPET)*, 2024, pp. 1748–1753.

[14] M. C. Silva Filho, R. L. Oliveira, C. C. Monteiro, P. R. M. Inácio, and M. M. Freire, "Cloudsim plus: A cloud computing simulation framework pursuing software engineering principles for improved modularity, extensibility and correctness," in *2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*, 2017, pp. 400–406.

[15] M. Tibaldi and C. Pilato, "Amethyst: Reducing data center emissions with dynamic autotuning and vm management," *IEEE Computer Architecture Letters*, vol. 24, no. 1, pp. 153–156, 2025.

[16] M. Tibaldi, S. Vinco, and C. Pilato, "Artemis repository," https://github.com/mattia-tibaldi/artemis.

[17] S. M. M. Amin, N. H. M. S. H. Lipu, S.Urooj, and A. Akter, "Development of a PV/battery micro-grid for a data center in Bangladesh: Resilience and sustainability analysis," vol. 15, 2023.

[18] B. Kim and I. Kim, "A case study of stand-alone hybrid power systems for a data center using HOMER and DIgSILENT," *Energy Reports*, vol. 9, pp. 1136–1143, 2023.

[19] S. Vinco, A. Sassone, F. Fummi, E. Macii, and M. Poncino, "An open-source framework for formal specification and simulation of electrical energy systems," in *Proc. of ISLPED*, 2014, p. 287–290.

[20] S. Dubey, J. N. Sarvaiya, and B. Seshadri, "Temperature dependent photovoltaic (PV) efficiency and its effect on PV production in the world – a review," *Energy Procedia*, vol. 33, pp. 311–321, 2013, pV Asia Pacific Conference 2012. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1876610213000829

[21] C. Kim, M.-C. Dinh, H.-J. Sung, K.-H. Kim, J.-H. Choi, L. Graber, I.-K. Yu, and M. Park, "Design, implementation, and evaluation of an output prediction model of the 10 mw floating offshore wind turbine for a digital twin," *Energies*, vol. 15, no. 17, 2022. [Online]. Available: https://www.mdpi.com/1996-1073/15/17/6329

[22] C. Liu, H. Wang, J. Ding, and C. Zhen, "An overview of modelling and simulation of thermal power plant," in *The 2011 International Conference on Advanced Mechatronic Systems*, 2011, pp. 86–91.

[23] "Determining the baseline efficiency of thermal or electric energy generation systems," https://cdm.unfccc.int/methodologies/PAmethodologies/tools/am-tool-09-v3.0.pdf, United Nation Framework Convention on Climate Change, Tech. Rep., 2020.

[24] M. Tomasov, M. Kajanova, P. Bracinik, and D. Motyka, "Overview of battery models for sustainable power and transport applications," *Transportation Research Procedia*, vol. 40, pp. 548–555, 2019, tRANSCOM 2019 13th International Scientific Conference on Sustainable, Modern and Safe Transport. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S235214651930242X

[25] "Nextsmartsaver - unlock cheaper electricity between 7pm-2am every day," https://www.eonnext.com/tariffs/next-smart-saver, 2025.

[26] National Center for environmental information (NOAA). (2018) Climate data online: Dataset discovery. [Online]. Available: https://www.ncei.noaa.gov/cdo-web/datasets

[27] K. Park and V. S. Pai, "CoMon: a mostly-scalable monitoring system for PlanetLab," *SIGOPS Oper. Syst. Rev.*, vol. 40, no. 1, p. 65–74, Jan. 2006. [Online]. Available: https://doi.org/10.1145/1113361.1113374

[28] C. Li, C. An, F. Yang, and X. Zeng, "Espsim:an efficient scalable power grid simulator based on parallel algebraic multigrid," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 28, no. 1, Dec. 2022. [Online]. Available: https://doi.org/10.1145/3529533

[29] T. Rendel, C. Rathke, T. Breithaupt, and L. Hofmann, "Integrated grid and power market simulation," in *2012 IEEE Power and Energy Society General Meeting*, 2012, pp. 1–7.

[30] S. N. Bhaskara and B. H. Chowdhury, "Microgrids — a review of modeling, control, protection, simulation and future potential," in *IEEE Power and Energy Society General Meeting*, 2012.

[31] A. Alzahrani, M. Ferdowsi, P. Shamsi, and C. H. Dagli, "Modeling and simulation of microgrid," *Procedia Computer Science*, vol. 114, pp. 392–400, 2017.

[32] M. A. Jirdehi, V. S. Tabar, S. Ghassemzadeh, and S. Tohidi, "Different aspects of microgrid management: A comprehensive review," *Journal of Energy Storage*, vol. 30, 2020.

[33] G. Shahgholian, "A brief review on microgrids: Operation, applications, modeling, and control," *International Transactions on Electrical Energy Systems*, vol. 31, no. 6, 2021.

[34] P. Wiesner, I. Behnke, and O. Kao, "A testbed for carbon-aware applications and systems," 2023.

[35] R. Andreoli, J. Zhao, T. Cucinotta, and R. Buyya, "CloudSim 7G: An integrated toolkit for modeling and simulation of future generation cloud computing environments," 2025. [Online]. Available: https://arxiv.org/abs/2408.13386