

Federated Learning-Based Digital Twin Framework for Decentralized Smart City Systems

¹Nishanth Hanchihalli Giriyappa

¹School of Electrical and Electronic Engineering, university college Dublin, Ireland

Email id: nishanth.hanchihalligiriya@ucdconnect.ie

Abstract *The rapid proliferation of Internet of Things (IoT) technologies has severely amplified the energy consumption bottlenecks of decentralized sensor networks. This paper presents a comprehensive review of energy-efficient Digital Twin (DT) architectures designed to mitigate the computational and transmission burdens placed on resource-constrained edge devices. By shifting from monolithic cloud structures to multi-layered, edge-centric deployments, DTs provide a synchronized virtual environment to intelligently offload complex tasks. We evaluate state-of-the-art AI-driven mechanisms, including Reinforcement Learning (D3QN) for adaptive task offloading and ensemble regression models (CatBoost) for proactive energy load forecasting. Furthermore, this review analyzes practical implementations in smart environments, demonstrating how non-intrusive occupancy monitoring and dynamic Building Information Modeling (BIM) can yield significant quantitative energy reductions, such as a 79% decrease in intelligent lighting costs. Ultimately, this framework establishes a pathway for transforming reactive IoT networks into proactive, energy-optimized ecosystems.*

Keywords Digital Twin, Energy Efficiency, Internet of Things (IoT), Edge Computing, Predictive Modeling, Task Offloading, Smart City.

I. Introduction

The transition towards intelligent environments—encompassing smart manufacturing, residential automation, and decentralized energy grids—has been heavily accelerated by the rapid deployment of Internet of Things (IoT) technologies. As these networks scale to incorporate millions of interconnected sensors, actuators, and edge devices, they generate immense volumes of real-time data crucial for continuous monitoring and automated decision-making. However, this unprecedented scale brings forth a critical operational challenge: managing the immense energy consumption required to sustain continuous sensing, data transmission, and local computation. To address this, the integration of Digital Twin (DT) technology has emerged as a highly effective architectural paradigm.[1][3] By creating a synchronized virtual replica of the physical network, DTs provide an environment to offload heavy computation, predict system behaviors, and optimize resource allocation dynamically, thereby extending the lifespan of resource-constrained IoT devices.[2]

While the deployment of large-scale IoT networks facilitates granular control over physical environments, it introduces severe operational constraints at the edge. In traditional architectures, physical nodes are tasked with continuous environmental sensing, localized data process-

ing, and high-frequency state transmissions to centralized servers. For battery-powered or resource-constrained devices, attempting to execute computationally intensive algorithms locally or maintaining continuous data synchronization rapidly depletes energy reserves. Consequently, achieving high energy efficiency—without sacrificing the computational intelligence required for complex tasks like occupancy detection or energy trading—has emerged as a primary bottleneck preventing the sustainable scaling of IoT networks.[4][5][6]

To mitigate these constraints, the DT paradigm represents a transformative architectural shift. Operating synchronously within edge or cloud environments, a DT serves as a dynamic, high-fidelity virtual representation of the physical IoT entities. By integrating a DT, heavy computational burdens, such as predictive modeling, system optimization, and complex task orchestration, are offloaded from the physical nodes. Rather than requiring devices to continuously process data or stream raw variables, the DT leverages advanced data-driven models to predict system states and simulate environmental dynamics. This architectural decoupling enables intelligent resource allocation, adaptive task offloading, and optimized energy routing, allowing the physical network to minimize unnecessary computation and significantly reduce overall power consumption.[7]

While the conceptual advantages of DTs are widely recognized, their specific architectural implementations for energy optimization in IoT networks vary significantly across different applications. This review synthesizes recent literature to provide a comprehensive analysis of state-of-the-art, energy-efficient DT architectures. The primary contributions of this paper are as follows:

- We evaluate multi-layered and edge-centric DT frameworks designed to decentralize data processing and reduce transmission latency.
- We analyze AI-driven mechanisms for intelligent task offloading and computational overhead reduction in resource-constrained networks.
- We review predictive modeling techniques utilized within DTs to forecast energy behavior, manage energy storage, and optimize real-time physical operations.
- We explore practical applications of these architectures in smart environments, including non-intrusive occupancy monitoring and dynamic building information modeling (BIM).

The remainder of this paper is organized as follows: Section II outlines proposed multi-layered DT architectures for IoT integration. Section III discusses AI-driven resource and task offloading strategies. Section IV reviews predictive modeling approaches for energy optimization. Section V explores practical application scenarios in smart environments, and Section VI concludes the paper.

II. Multi-Layered DT Architectures for Energy-Efficient IoT

To successfully mitigate the energy bottlenecks inherent in large-scale IoT networks, the integration of Digital Twins must be systematically structured. Deploying a DT as a monolithic, cloud-based entity often exacerbates latency and communication energy costs. Consequently, recent literature emphasizes the adoption of multi-layered, edge-centric architectural frameworks that distribute computational loads and prioritize localized data processing.

A. The Six-Layer Conceptual Framework

A foundational approach to energy-efficient DT integration is the deployment of a highly modular, multi-tier architecture. Recent research proposes a comprehensive six-layer conceptual framework designed specifically to support the integration of real-time data, AI models, and control logic while minimizing energy overhead. This architecture comprises the following interconnected tiers:

1. **Physical/Edge Layer:** Encompasses the hardware IoT devices, actuators, and sensor networks responsible for initial data acquisition.

2. **Communication Layer:** Manages the energy-efficient transmission protocols routing data between the physical environment and the virtual counterpart.
3. **Data Management Layer:** Handles the aggregation, filtering, and localized storage of telemetry data, ensuring that only essential, pre-processed information is forwarded to computationally intensive models.
4. **AI and Modeling Layer:** The core of the DT, housing the data-driven algorithms and predictive models (e.g., Digital Shadows) that simulate system states.
5. **User Interface (UI) Layer:** Provides visualization and actionable recommender systems to end-users or building managers.
6. **Security and Governance Layer:** Cross-cuts all tiers to ensure data integrity and manage access without imposing restrictive, energy-draining cryptographic overhead on edge nodes.

By distinctly separating the Data Management and AI layers from the physical hardware, this framework allows predictive models to operate independently of continuous sensor polling. Furthermore, the utilization of open-source, event-driven frameworks—such as Node-RED—facilitates the lightweight deployment of these architectural layers, enabling seamless data flow and algorithmic execution with minimal computational friction.

B. Edge-Centric Deployment and Virtualization

To further optimize energy consumption, contemporary architectures heavily favor edge-centric DT deployments over traditional cloud-reliant models. By instantiating the DT directly on edge devices (such as a local Raspberry Pi gateway), networks can drastically reduce the transmission energy and latency associated with continuous cloud synchronization.

A prominent example is the integration of DTs within edge-based home automation platforms, such as Home-Assistant. In this architecture, localized data storage and real-time processing occur proximal to the data source. This proximity allows the DT to execute rapid, energy-saving control decisions (e.g., adaptive HVAC or lighting adjustments) without requiring round-trip cloud communication.[8][9]

Moreover, edge-centric DT architectures facilitate the implementation of "virtual sensors." Rather than deploying power-hungry physical sensors for every environmental variable, the DT synthesizes available heterogeneous data to infer missing parameters. This virtualization enables the network to maintain a comprehensive understanding of the physical environment—such as inferring room occupancy or thermal dynamics—while significantly reducing the hardware footprint and the cumulative energy consumed by physical sensing operations.

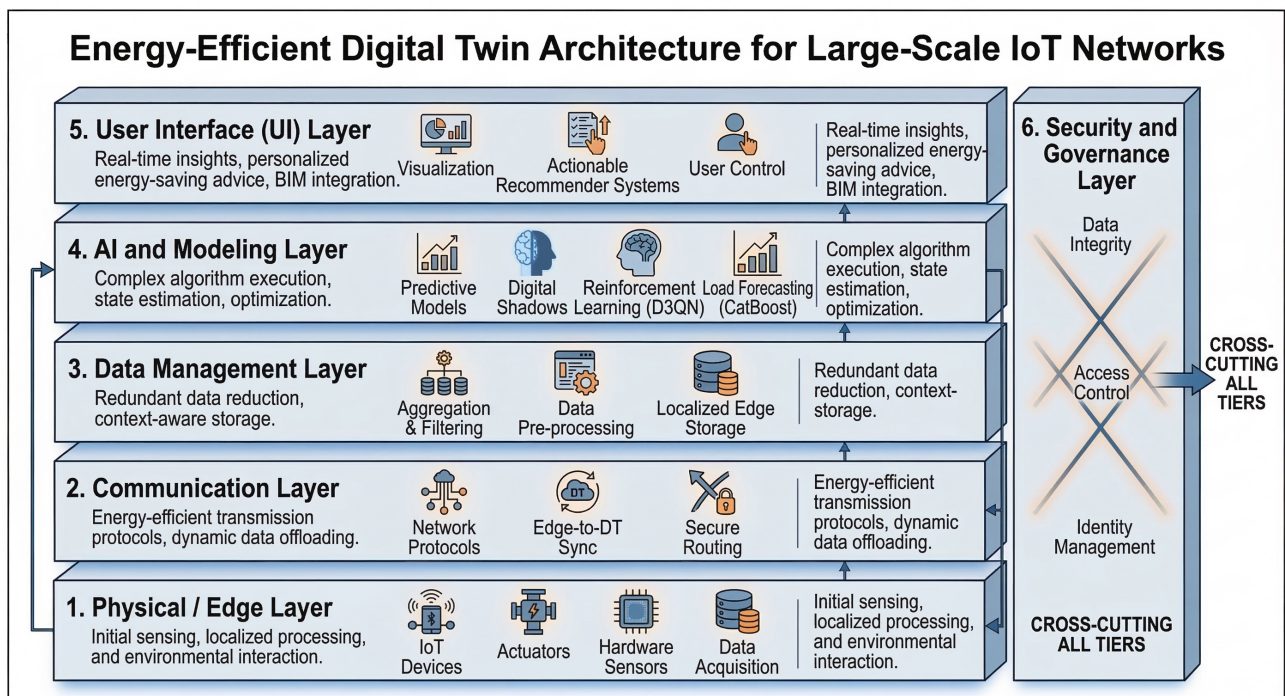


Figure 1: The proposed six-layer conceptual architecture for energy-efficient Digital Twin integration in IoT networks.

III. AI-Driven Resource and Task Offloading

As Digital Twin architectures migrate towards the edge to reduce transmission latency, managing the localized computational load becomes paramount. Edge devices and physical sensor nodes often possess limited battery capacities and processing capabilities. Consequently, requiring these nodes to continuously execute complex algorithms or process high-dimensional data streams locally leads to rapid energy depletion. To resolve this, contemporary DT frameworks increasingly incorporate Artificial Intelligence (AI) to intelligently orchestrate resource allocation and dynamically offload tasks.

A. Reinforcement Learning for Adaptive Offloading

A highly effective method for managing distributed computational tasks involves integrating Reinforcement Learning (RL) within the DT layer. Recent advancements propose the utilization of advanced RL models, specifically the Dueling Double Deep Q-Network (D3QN), to govern offloading decisions in IoT networks.[10]

In this architecture, the DT plays a critical preparatory role by generating predictive analytics regarding the physical node's state. By employing techniques such as linear regression to forecast battery drain and moving averages to estimate impending CPU loads, the DT provides a high-fidelity simulation of the node's energy trajectory.[11] This predictive data is then fed to the D3QN agent, which dy-

namically decides whether an incoming computational task should be executed locally on the resource-constrained sensor or offloaded to a more powerful edge server. By continually learning the optimal trade-off between communication energy costs (for offloading) and computational energy costs (for local execution), this DT-driven RL approach has been demonstrated to successfully maintain local sensor battery levels above 85% during sustained operations, vastly extending the operational lifespan of the network.[12][13]

B. Algorithmic Optimization for Heavy Workloads

Beyond deciding *where* to process a task, AI-driven DTs are also utilized to determine *how much* data requires processing, specifically for inherently intensive tasks such as computer vision. In smart building environments, utilizing complex object detection models (e.g., YOLOv4) for pedestrian tracking imposes massive computational overhead on edge devices.

To mitigate the associated energy drain, DT architectures implement intelligent filtering mechanisms prior to algorithmic execution. By integrating Ambient Brightness Perception and Key-Frame Similarity assessments, the system evaluates the visual data stream to discard redundant or uninformative frames before they reach the heavy processing pipeline. For instance, if the DT detects no significant pixel variation between consecutive frames, or if the ambient lighting renders the frame unanalyzable, the processing task is bypassed. This selective processing approach drastically

reduces unnecessary hardware activation, achieving approximately a 35% reduction in GPU utilization and a 64% reduction in CPU utilization, thereby converting a traditionally power-hungry application into an energy-efficient edge deployment.[14]-[16]

IV. Predictive Modeling for Energy Optimization

A fundamental advantage of the Digital Twin architecture is its capacity to transition IoT networks from reactive energy management to proactive energy optimization. By leveraging historical data and real-time telemetry, the DT employs advanced predictive modeling to forecast energy demands, simulate system dynamics, and orchestrate energy trading before physical execution is required. This predictive capability is essential for minimizing energy waste across both industrial and residential environments.

A. Short-Term Load and Generation Forecasting

Accurate forecasting of energy consumption allows the DT to implement anticipatory energy-saving strategies, such as load shifting or proactive HVAC modulation. Recent frameworks highlight the efficacy of ensemble regression models in these scenarios. For instance, the deployment of the CatBoost algorithm within building automation DTs has demonstrated exceptional predictive accuracy ($R^2 \geq 0.92$) for short-term energy consumption. By accurately anticipating thermal dynamics and energy loads, the DT can optimize the activation cycles of high-power appliances, significantly reducing peak demand.

Furthermore, in decentralized energy environments such as peer-to-peer (P2P) nano-grids, prediction modules utilize deep learning architectures like Gated Recurrent Units (GRUs). Within the DT, GRUs are effectively employed to forecast both impending energy loads and Photovoltaic (PV) generation capacity. This dual-forecasting mechanism ensures that the system possesses a reliable estimation of future energy states, which is a prerequisite for executing efficient energy trading and storage operations.

B. Dynamic Energy Behavior Modeling

In complex industrial IoT settings, such as manufacturing systems, traditional state-based energy modeling often falls short. These conventional models typically assume constant power consumption during operational states, which severely distorts the actual, highly variable energy usage of industrial machinery.

To achieve high-fidelity energy optimization, recent literature proposes the integration of Data-Driven Hybrid Petri-Nets (DDHPN) within the DT. When combined with advanced machine learning techniques like the Gaussian Kernel Extreme Learning Machine (KELM), this approach allows for the accurate, real-time fitting of instantaneous firing speeds across continuous energy transitions. By dynam-

ically modeling the exact energy behavior of the manufacturing process, the DT provides a precise virtual space where engineers can simulate and refine energy-saving protocols without disrupting the physical production line.

C. Optimizing Energy Storage and Trading

Predictive modeling within the DT is not solely utilized for load reduction; it is also critical for optimizing the distribution of surplus energy. In large-scale, decentralized IoT networks featuring Energy Storage Systems (ESS), managing the charge and discharge cycles efficiently is vital for overall network sustainability.

Architectures have successfully integrated heuristic algorithms, specifically Particle Swarm Optimization (PSO), alongside IoT task orchestrators within the DT layer. Using the load and generation forecasts provided by the predictive models (e.g., GRUs), the PSO algorithm dynamically calculates the optimal energy routing strategy. This ensures that surplus energy is efficiently allocated to the ESS or traded within the P2P network, thereby minimizing energy trading costs and maximizing the utilization of renewable energy sources across the physical grid.

V. Application Scenarios in Smart Environments

The theoretical frameworks and predictive models of energy-efficient Digital Twins find their most immediate and impactful applications within smart environments, particularly in residential automation and commercial building management. By synthesizing real-time data, DTs enable sophisticated, energy-saving interventions that operate seamlessly without disrupting user comfort.

A. Non-Intrusive Occupancy Monitoring and Recommender Systems

A primary driver of energy waste in intelligent buildings is the continuous operation of HVAC and lighting systems in unoccupied spaces. Traditional occupancy detection often relies on continuous camera surveillance, which is computationally expensive, power-intensive, and raises privacy concerns. To resolve this, edge-centric DT architectures leverage non-intrusive environmental data to infer human presence.

By aggregating low-power ambient telemetry—such as temperature, humidity, ambient light, carbon dioxide (CO_2) levels, and basic motion events—machine learning models housed within the DT can predict room occupancy with exceptional precision. Recent implementations demonstrate that such multi-variable data fusion achieves an occupancy detection accuracy of up to 95.12%. Furthermore, this highly accurate state estimation feeds directly into DT-integrated recommender systems. Instead of strictly automated, rigid control, these systems analyze occupancy trends and environmental conditions to push personalized,

Table 1: Summary of AI Algorithms and Quantitative Energy Metrics in DT-IoT Frameworks

Application Area	AI/ML Algorithm	Primary Energy Benefit	Quantitative Metric
Task Offloading	D3QN (Reinforcement Learning)	Prolongs edge node battery	Maintains battery > 85%
Computer Vision Filtering	YOLOv4 with Key-Frame Similarity	Reduces hardware activation	64% CPU / 35% GPU utilization drop
Short-Term Load Forecasting	CatBoost (Ensemble Regression)	Proactive HVAC load shifting	High predictive accuracy ($R^2 > 0.92$)
Energy Trading & Storage	GRU + Particle Swarm Optimization	Minimizes energy trading costs	Optimizes ESS charge/discharge

actionable energy-saving recommendations to users, such as advising the utilization of natural ventilation during optimal external weather conditions.

B. Dynamic BIM and Adaptive Lighting Control

In commercial infrastructure, lighting constitutes a massive portion of total energy expenditure. Traditional intelligent lighting systems generally rely on isolated, localized sensor control without integrating multi-source data. Contemporary DT architectures solve this by combining computer vision with dynamic Building Information Modeling (BIM).

In this application, the DT serves as a visualized operation and maintenance platform. It fuses real-time surveillance video streams with the spatial data provided by the BIM. By employing the optimized, energy-filtered computer vision algorithms discussed in Section III, the DT accurately maps pedestrian presence and measures ambient illuminance within specific geometric zones. The DT then adaptively adjusts the Pulse Width Modulation (PWM) of dimmable LED lighting systems based on exact, real-time spatial requirements rather than pre-programmed schedules. This holistic, data-driven synchronization between the physical lighting infrastructure and the dynamic BIM virtual space has been shown to achieve a remarkable 79% reduction in lighting electricity costs.




Figure 2: Data flow for dynamic BIM-integrated adaptive lighting control.

C. Open Challenges and Future Directions

Despite the significant energy reductions facilitated by DT integration, several critical challenges remain. First, the energy paradox of the Digital Twin itself must be addressed; running complex, high-fidelity AI models at the edge introduces its own computational energy costs, which must

be strictly balanced against the energy saved at the physical nodes. Future research must prioritize the development of ultra-lightweight, quantized machine learning models tailored specifically for edge-DT environments. Secondly, interoperability remains a massive hurdle in decentralized smart city systems. Large-scale IoT networks consist of highly heterogeneous devices utilizing disparate communication protocols. Establishing standardized, low-overhead semantic frameworks that allow seamless data synchronization between legacy physical devices and modern DT architectures is essential for future scalability.

VI. Conclusion

As the Internet of Things scales to encompass millions of distributed nodes across smart manufacturing, decentralized grids, and intelligent buildings, the energy constraints of physical devices present a critical barrier to sustainable operation. This review has demonstrated that the integration of Digital Twin technology provides a robust architectural solution to the IoT energy bottleneck. By transitioning away from monolithic cloud structures toward multi-layered, edge-centric frameworks, networks can significantly reduce transmission latency and localize data management.

Crucially, the DT paradigm enables the offloading of heavy computational tasks. Through the deployment of AI-driven reinforcement learning agents, such as D3QN, and the implementation of intelligent data filtering, resource-constrained physical nodes can drastically reduce GPU and CPU utilization, extending their operational lifespans. Furthermore, the integration of advanced predictive models—including ensemble regression, Gated Recurrent Units, and dynamic Petri-nets—allows the DT to accurately forecast energy loads, simulate physical states, and optimize energy trading without requiring continuous, power-draining physical synchronization. Ultimately, the deployment of energy-efficient Digital Twins transforms IoT networks from reactive, sensor-heavy deployments into proactive, intelligent ecosystems capable of maximizing computational power while minimizing their physical energy footprint.

References

[1] L. Belussi, B. Barozzi, A. Bellazzi, L. Danza, A. Devitofrancesco, C. Fanciulli, M. Ghellere, G. Guazzi, I. Meroni, F. Salamone, et al., “A review of performance

- of zero energy buildings and energy efficiency solutions,” *Journal of Building Engineering*, vol. 25, 2019, Art. no. 100772.
- [2] S. Chen, G. Zhang, X. Xia, S. Setunge, and L. Shi, “A review of internal and external influencing factors on energy efficiency design of buildings,” *Energy and Buildings*, vol. 216, 2020, Art. no. 109944.
- [3] M. U. Mehmood, D. Chun, H. Han, G. Jeon, and K. Chen, “A review of the applications of artificial intelligence and big data to buildings for energy efficiency and a comfortable indoor living environment,” *Energy and Buildings*, vol. 202, 2019, Art. no. 109383.
- [4] Z. Hu, W. Yao, Z. Shi, H. Shuai, W. Gan, and J. Wen, “Intelligent and rapid event-based load shedding pre-determination for large-scale power systems: Knowledge-enhanced parallel branching dueling Q-network approach,” *Applied Energy*, vol. 347, 2023, Art. no. 121468.
- [5] Y. Jiang and T. H. Ortmeier, “Propagation-based network partitioning strategies for parallel power system restoration with variable renewable generation resources,” *IEEE Access*, vol. 9, pp. 144965–144975, 2021.
- [6] F. Tao and M. Zhang, “Digital twin shop-floor: A new shop-floor paradigm towards smart manufacturing,” *IEEE Access*, vol. 5, pp. 20418–20427, 2017.
- [7] B. Yang, Z. Lv, and F. Wang, “Digital twins for intelligent green buildings,” *Buildings*, vol. 12, no. 6, Art. no. 856, 2022.
- [8] W. Huang, Y. Zhang, and W. Zeng, “Development and application of digital twin technology for integrated regional energy systems in smart cities,” *Sustainable Computing: Informatics and Systems*, vol. 36, 2022, Art. no. 100781.
- [9] A. Rasheed, O. San, and T. Kvamsdal, “Digital twin: Values, challenges and enablers from a modeling perspective,” *IEEE Access*, vol. 8, pp. 21980–22012, 2020.
- [10] R. Bortolini, R. Rodrigues, H. Alavi, L. F. D. Vecchia, and N. Forcada, “Digital twins’ applications for building energy efficiency: A review,” *Energies*, vol. 15, no. 19, Art. no. 7002, 2022.
- [11] W. Yu, P. Patros, B. Young, E. Klinac, and T. G. Walmesley, “Energy digital twin technology for industrial energy management: Classification, challenges and future,” *Renewable and Sustainable Energy Reviews*, vol. 161, 2022, Art. no. 112407.
- [12] S. Mihai, M. Yaqoob, D. V. Hung, W. Davis, P. Towakel, M. Raza, M. Karamanoglu, B. Barn, D. Shetve, and R. V. Prasad, “Digital twins: A survey on enabling technologies, challenges, trends and future prospects,” *IEEE Communications Surveys & Tutorials*, 2022.
- [13] Z. Tu, L. Qiao, R. Nowak, H. Lv, and Z. Lv, “Digital twins-based automated pilot for energy efficiency assessment of intelligent transportation infrastructure,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 11, pp. 22320–22330, 2022.
- [14] M. Alford, I. Udugama, W. Yu, and B. Young, “Flexible digital twins from commercial off-the-shelf software solutions: A driver for energy efficiency and decarbonisation in process industries?” *Chemical Product and Process Modeling*, vol. 17, no. 4, pp. 395–407, 2022.
- [15] Y. Tan, P. Chen, W. Shou, and A. M. Sadick, “Digital twin-driven approach to improving energy efficiency of indoor lighting based on computer vision and dynamic BIM,” *Energy and Buildings*, vol. 270, 2022, Art. no. 112271.
- [16] A. Saad, S. Faddel, and O. Mohammed, “IoT-based digital twin for energy cyber-physical systems: Design and implementation,” *Energies*, vol. 13, no. 18, Art. no. 4762, 2020.