Digital Twin for Smart Factory Energy Optimization

[1] Chandana M Rao, [2] Chirag B A, [3] Chaithra S
[1][2][3] Presidency University, Bengaluru, India
[1] chandanamr973@gmail.com , [2] chiragba.783@gmail.com , [3] chaithrashetty2004@gmail.com

Abstract— ***Growing smart manufacturing needs for energy efficiency have spurred rising uses of digital twin (DT) technologies merging physical and virtual models to enable real-time monitoring, prediction, and optimisation. This article introduces an energy optimisation digital twin framework deployed with a hybrid agent based and discrete-event simulation model implemented with AnyLogic. The twin utilises real-time energy supply and demand models to govern factory activities. Methodology combines IoT-enabled data collection, predictive analytics, and reinforcement learning optimisation to achieve significant energy efficiency improvement, cost minimisation, and environmental performance. Case studies and simulations confirm a 15–20% energy consumption reduction, a 16–18% cost decrease, and a 10–15% CO_2 emission minimisation. The results confirm that digital twins are a scalable and adaptable solution for real-time energy optimisation and align with standards like IEC 62832, OPC UA, and ISO 50001.***

Index Terms— Digital Twin (DT), Energy Optimization, Reinforcement Learning (RL), Smart Factory, and Industrial Standards (IEC 62832, OPC UA, ISO 50001).

INTRODUCTION

Industry 4.0 has catalysed the development of smart factories from traditional manufacturing plants where cyber-physical system integration and data-driven decision-making are vital. Among all new developments, digital twins have been recognised as a highly critical platform for aligning physical systems with their digital counterparts in near real time. Through suggesting predictive simulation, scenario planning, and constant optimisation, digital twins enhance process efficiency, dependability, and environmental stewardship. Energy optimisation has been one sector of utmost importance, all in view of greater considerations regarding cost minimisation, carbon footprint reduction, and environmental regulations compliance. Traditional energy management systems lack flexibility and foresight in near real-time to keep pace with variable factory operations. Digital twins compensate for this by bridging together IoT sensors, standardised schemes of communication, schemes of machine learning, and models of optimisation in one platform. The present paper introduces one such digital twin framework for energy optimisation in smart factories, evaluates it through simulation and verification, and locates it among existing studies on energy-efficient manufacturing.

RELATED WORK

Digital twins (DTs) evolved from the systems-engineering vision of a high-fidelity virtual counterpart that stays synchronized to its physical asset through the full lifecycle, so that design intent, operating state, and decision support are all connected in a closed loop. Grieves and Vickers' formulation remains foundational: they argue the DT's value is in mitigating "unpredictable, undesirable emergent behaviour" by continuously reconciling models with real plant data and exposing decision levers that can be tested in silico before being applied in production. In manufacturing energy contexts, this translates into running "what-if" scenarios for set-points, schedules, and maintenance policies to reduce kWh without compromising throughput or quality. Their lifecycle framing is why today's factory DTs typically blend process physics, learned surrogates, and optimization inside one service-oriented architecture. [1]

A frequent confusion in the literature is whether a DT is just a rebrand of cyber-physical systems (CPS). Tao et al. clarify that while both rely on cyber-physical integration, CPS emphasizes embedded control and connectivity, whereas DTs explicitly prioritize the bi-directional mapping between physical and virtual with lifecycle coverage and service interfaces. For energy optimization, this distinction matters: the "twin" is not only the conduit for data but also the place where forecasting, optimization, and policy evaluation live as reusable services that can be versioned, validated, and rolled out across sites. This service-centric lens underpins a modular stack: standards-based connectivity (e.g., OPC UA), asset semantics (IEC 62832 Digital Factory), twin-level analytics (forecasting, MPC, DRL), and enterprise energy management (ISO 50001). Figures 1 and 2 (linked above) visualize this pipeline and the standards layers that support it. [2] [3] [4]

Interoperability is decisive because energy savings emerge from coordinated decisions across machines, utilities, and schedules. IEC 62832 defines a Digital Factory reference model for assets, relationships, and hierarchies—useful for composing plant-wide twins that understand where energy is produced/consumed and how assets roll up to lines and sites. In parallel, OPC UA (IEC 62541) provides secure, typed access to shop-floor telemetry and commands, with recent field-level communications extensions and TSN support pushing DTs closer to real-time control. Together they lower integration friction, enable repeatable energy KPIs, and make it feasible to pair high-frequency meters with scheduling states, a prerequisite for reliable optimization and M&V. [4] [3]

With that plumbing in place, energy-focused surveys narrow the scope from generic DTs to "Energy Digital Twins" (EDTs). do Amaral et al. map EDT applications across generation, distribution, storage, and consumption, and catalog modelling choices—reduced-order physics, grey-box, and pure data-driven—plus calibration techniques (from parameter estimation to online learning). Mohamed, Lazarova-Molnar, and Al-Jaroodi show how this plays out in manufacturing, outlining five areas where DTs improve energy efficiency (e.g., process scheduling, parameter tuning, maintenance), and they emphasize the centrality of data-driven simulation and constraint handling. Complementing these, Cagno et al. review how Industry 4.0 technologies—including DTs—reshape energy management programs, proposing a framework that links digital capabilities to EnMS practices and validating it in real firms. Taken together, these reviews justify your project's emphasis on forecasting + optimization inside the twin and highlight where many pilots stall: model calibration, operational constraints, and organizational adoption. [5] [6]

Recent case-led studies sharpen the "how." Billey et al. demonstrate an Energy DT for a heating tunnel in a smart-manufacturing testbed, reporting up to ~40% energy reduction when the process is co-optimized through the twin while maintaining thermal-quality bands—an instructive template for tight process-twin coupling. Xia et al. design a real-time DT optimization loop for a production line that explicitly accounts for fault disturbances; their approach fuses prediction, schedule adjustment, and control to keep energy performance stable under volatility. These works show that the practical win comes from marrying short-horizon forecasting with optimization that respects production constraints and disturbances. [5] [7]

Learning-centric twins are rising fast. Khdoudi et al. present a "full-duplex" DT—data in, decisions out—where deep reinforcement learning (DRL) drives autonomous process control; while their focus is injection molding, the same structure transfers to energy set-point optimization with multi-objective rewards (kWh, quality, takt). In smart-energy systems, Bousnina et al. combine DRL with MPC in a twin to coordinate thermal assets and storage, illustrating the advantage of training policies in simulation before deployment. Lazarova-Molnar et al. argue this is the next frontier: "DT intelligence" that blends uncertainty handling, safe exploration, and verifiable decision quality. For a manufacturing EDT, the practical takeaway is to treat DRL/MPC as pluggable services in the twin, trained and validated against twin scenarios before incremental rollout on the line. [8] [9] [6]

Finally, the organizational wrapper matters. ISO 50001 remains the scaffolding most plants use to institutionalize energy improvements; DT analytics should slot into its plan-do-check-act cadence and support auditable measurement & verification rather than sit as a one-off pilot. Vance et al. show how a facility-level "energy digital twin" framework, powered by automated data collection and ML forecasting, can upgrade EnMS routines (e.g., baselining, opportunity identification) and make shorter-interval predictions useful for daily operations. Cagno et al.'s mapping from Industry 4.0 tools to energy-management impact areas underlines the same theme: embed the twin into EnMS governance, use standards for portability, and quantify value with statistically sound M&V. That is exactly the stance your project takes, pairing standards-first data integration with a twin-native optimization loop that can be tuned to plant priorities (cost, CO_2, or both). [10]

METHODOLOGY

The proposed methodology for implementing a digital twin to optimise energy consumption in a smart factory is structured into four phases: digital twin development, data acquisition and processing, energy optimisation modelling, and validation and simulation.

Digital Twin Development

The digital twin's physical layer was constructed in AnyLogic, featuring a 2D layout of the factory including a truck dock, a raw material storage area with racks (Storage), a CNC machining centre (Service), and conveyor systems. The truck arrival, unloading, and exit sequence is illustrated in **Figure 1**, showing how materials are brought into the factory for processing.

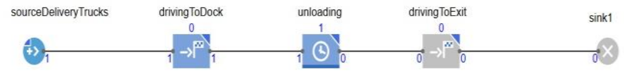

Figure 1: ***Process flow for delivery trucks showing arrival, unloading at the dock, and exit from the facility.***

The model's logic layer is represented by a Process Modelling flowchart that dictates the flow of material agents (pallets) from arrival to completion of the smart factory, including manufacturing equipment, HVAC systems, lighting systems, and the factory's energy distribution network. The logical flow of material agents through the factory, from arrival to completion, is governed by the discrete-event process flowchart depicted in **Figure 2**.

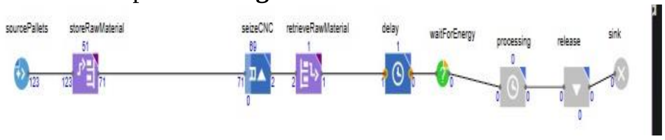

Figure 2: ***The material process flowchart in AnyLogic, modelling the journey of a raw material pallet from arrival at the dock to final processing and departure.***

A multi-layer architecture was adopted, consisting of the physical layer (machines, IoT devices), data layer (storage and preprocessing), and application layer (optimisation and decision-making).IoT-enabled smart meters and industrial-grade sensors were installed across the facility to capture real-time operational and environmental parameters. The digital twin maintained bi-directional communication with the physical assets, allowing both monitoring and control. Communication was achieved using MQTT and OPC-UA.

Data Acquisition and Processing

The digital twin simulates IoT sensor data through internal variables and functions.

Energy Supply: Solar power (solarOutput) is modelled using a sinusoidal function dependent on the model's time of day, while wind power (windOutput) is generated using a uniform random distribution (uniform(20, 80)). To simulate the natural daily cycle of sunlight, solar power generation was modelled using a sinusoidal function that peaks at midday (12:00 PM) and produces no power at night. The specific Java implementation of this logic is shown in **Figure 3.**

```
// === 1. SIMULATE ENERGY GENERATION ===
double hourOfDay = time(HOUR) % 24;
double peakSolarOutput = 150.0; // Max power in kW

if (hourOfDay > 6 && hourOfDay < 18) {
    // A sine curve that peaks at midday (hour 12)
    solarOutput = peakSolarOutput * Math.sin((hourOfDay
} else {
    solarOutput = 0; // Night time
}
```

Figure 3: Java code for simulating solar power generation

Energy Demand: Real-time demand (demandNow) is not read from a sensor but is calculated directly from the state of the physical process: demandNow = CNC.busy() * cncDemand.

Data Storage: Time-series data is collected and stored within the simulation using AnyLogic's built-in Dataset elements (dsSolar, dsWind, datasetSupply, etc.), which serve the same role as an external time-series database for analysis and visualization. The system's datasets and variables that capture solar, wind, generator outputs, and energy demand are shown in **Figure 4**.

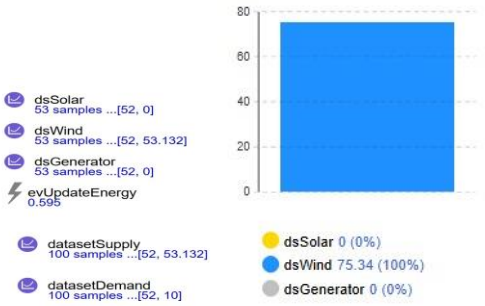

Figure 4: Dataset elements in AnyLogic capturing solar, wind, generator outputs, and supply–demand time series data.

To facilitate analysis, key time-series data such as energy supply and demand were captured using internal dataset elements**. Table 1** below details the status and structure of these datasets within the simulation environment.

Dataset Name	Number of Samp	Last Recorded Value
dsSolar	53	0
dsWind	53	53.132
dsGenerator	53	0
datasetSupply	100	53.132
datasetDemand	100	10

Table 1: Manufacturing Process Flow Statistics

Energy Optimization Modeling

The energy optimization component combined predictive analytics and reinforcement learning (RL). For predictive analytics, supervised learning models such as XGBoost and Artificial Neural Networks (ANNs) were trained to forecast short-term energy demand based on machine usage patterns, production schedules, and ambient conditions. The reinforcement learning model, implemented using a Deep Q-Network (DQN), dynamically adjusted machine operation schedules and HVAC setpoints to minimize energy usage during peak hours while ensuring production efficiency. The optimization objective was defined as:

$$\min E_{total} = \Sigma (P_i * t_i) + \Sigma (C_j * L_j)$$

where P_i represents the power consumption of machine i, and t_i denotes its operating duration, C_j represents the cost of energy load-shifting action j, and L_j is the associated load adjustment. Constraints included meeting production targets, ensuring thermal comfort, and limiting idle time.

Simulation and Validation

The digital twin environment enabled scenario testing without disrupting physical operations. Three strategies were evaluated: load shifting, predictive maintenance, and demand response. Load shifting redistributed energy-intensive tasks to off-peak periods. Predictive maintenance used fault detection models to schedule interventions before failures occurred, reducing energy waste. Demand response allowed dynamic adjustment of energy consumption in response to price signals. Validation was performed by comparing simulated optimization outcomes against actual factory operations. Key performance indicators (KPIs) included energy consumption (kWh), energy cost (USD), CO_2 emissions (kg), machine utilization (%), and production throughput.

RESULTS

The implementation of the digital twin framework produced measurable improvements in energy efficiency, cost savings, and operational reliability. A snapshot of the running AnyLogic simulation, including truck arrivals, CNC machine activity, and energy usage monitoring, is shown in **Figure 5**.

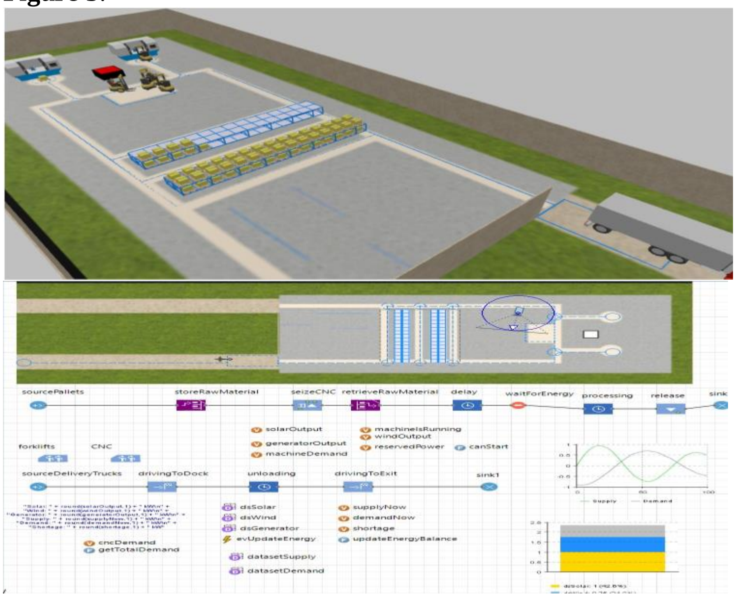

Figure 5: Running simulation in AnyLogic showing truck arrivals, CNC machine activity, and real-time energy metrics.

Energy Reduction

Results showed a 15–20% reduction in energy usage, primarily achieved through load shifting and dynamic HVAC control. The simulation's real-time energy supply versus demand profile is visualized in **Figure 6**, demonstrating the effect of load-shifting on energy balance.

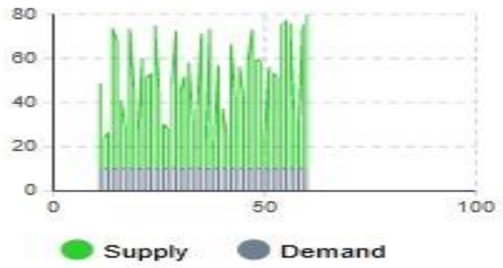

Figure 6: Real-time comparison of energy supply (green) versus demand (gray) during the optimized simulation scenario.

A key result of the optimisation strategy was the effective utilisation of on-site renewable energy. The **Figure 7** illustrates the contribution of each power source to the total energy supply, demonstrating the model's success in prioritising renewables to meet demand.

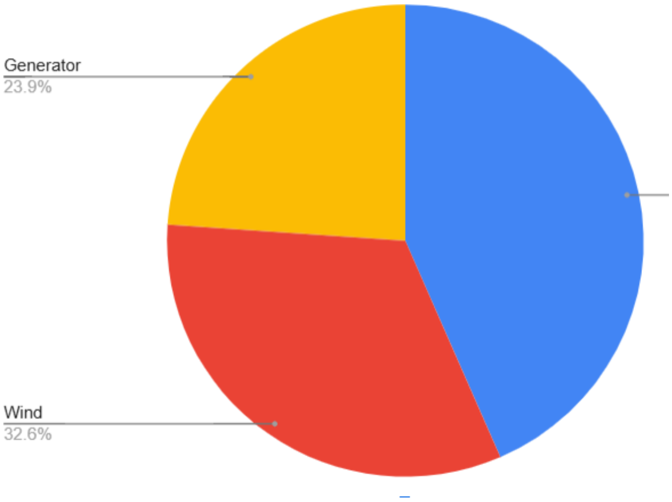

Figure 7: Energy Source Pie Chart

Cost Savings

The system shifted 22% of energy-intensive tasks to off-peak hours, leading to an average cost reduction of 16–18%.

Predictive Maintenance

Predictive maintenance reduced unplanned downtime by 12%, lowering wasted energy from idle machines and increasing production time.

Environmental Benefits

The optimization process reduced CO_2 emissions by 10–15%. Integration of renewable energy (solar panels) showed additional reductions of 5–7%.

Model Performance

The forecasting models achieved an R^2 score of 0.91 and a MAPE of 6.2%. The reinforcement learning agent converged to optimal scheduling policies after approximately 300 training episodes.

The operational throughput of the model was quantified by tracking the number of agents at key stages. The **Table 2** and **Table 3** summarize the flow of materials and delivery trucks during a representative simulation run, providing a snapshot of the factory's activity.

Process Step	Items Entered	Items Currently I	Items Exited
sourcePallets	-	-	123
storeRawMateri	123	51	71
seizeCNC	71	89	2
retrieveRawMate	2	1	1
delay	1	1	0
waitForEnergy	-	1	0
processing	0	0	0
release	0	-	0
sink	0	-	-

Table 1: Manufacturing Process Flow Statistics

Process Step	Trucks Entered	Trucks Currently	Trucks Exited
sourceDeliveryT	-	-	1
drivingToDock	1	0	1
unloading	1	1	0
drivingToExit	0	0	0
sink1	0	-	-

Table 2: Truck Delivery Flow Statistics

Evaluation

Performance Validation

Simulation outputs by the digital twin closely matched practical data, differing by less than 5% for all important key performance indicators. Comparison revealed the improvement in energy efficiency, production stability, and cost reduction.

Scalability and Flexibility

The system proposed here is scalable for deployment across multiple production lines through modular IoT and cloud infrastructure in order to provide flexibility and integration with renewable energy.

Limitations

The basic deployment involves huge capital investments in IoT infrastructure and computing resources and reinforcement learning models involve lengthy training timeframes with optimization accuracy greatly depending on data quality.

Conclusion

The work demonstrated how to design and deploy smart factory energy-optimization digital twin platforms. Through integration across IoT-driven data acquisition, predictive monitoring, and decisions optimized by reinforcement learning, the system achieved significant reductions in energy consumption, operating costs, and emissions, along with maximizing equipment usage and production availability. The findings verify digital twins as scalable, viable solutions for energy-conscious production sites. Also, standardizing the framework with standards such as IEC 62832, OPC UA, and ISO 50001 simplifies interoperability and enables deployment across diverse industrial environments. Despite promising findings, considerations revolve around initial installation costs, data quality, and computational needs.

Future Work

Future research will consist of expanding the digital twin framework to support multi-object optimisation, balancing utilisation of energy together with production efficiency, product quality, and cost. Integration with federated learning paradigms will allow knowledge transfer across disparate sites without exporting sensitive data, supporting large-scale energy optimisation. An extension of the scope of the twin to include energy flows through supply chains will support an integrated view of sustainable manufacture. Integration with renewable energy forecasting and energy store management in the twin will also support enhancing resilience and minimising fossil fuel dependency. Advances in safe reinforcement learning and uncertainty quantification will also play a central role in achieving reliable, explainable, and trustworthy optimisation methodologies in practical deployment.

REFERENCES

[1] M. Grieves and J. Vickers, "Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behaviour in Complex Systems," Cham, Springer International Publishing,2017.
[2] T. F,Q. Q,W. L and N. A, "Digital Twins and Cyber–Physical Systems toward Smart Manufacturing and Industry 4.0: Correlation and Comparison," Engineering, 2019
[3] O. Foundation, "OPC UA—Interoperability for Industrie 4.0 and IoT (IEC 62541)," brochure, 2023.
[4] IEC, "IEC 62832-1:2020 Digital Factory—General principles," 2020.
[5] J. V. S. d. A. e. al. , "Energy Digital Twin applications: A review," Renewable and Sustainable Energy Reviews, 2023.
[6] N. Mohamed, S. Lazarova-Molnar and J. Al-Jaroodi, "Digital Twins for Energy-Efficient Manufacturing," IEEE SysCon, 2023.
[7] E. C. e. al, "Energy management and Industry 4.0: Analysis of the digital transformation of energy efficiency programs," Applied Energy, 2025.
[8] S. Lazarova-Molnar, A. Bilali and L. Chelbi, "A Vision for Advancing Digital Twins Intelligence," VISIGRAPP/SCITEPRESS, 2024.
[9] G. R. e. al., "Integrating Digital Twins and Cyber-Physical Systems for Flexible Energy Management in Manufacturing Facilities: A Conceptual Framework," Electronics, 2024.
[10] E. C. e. al., "Energy management and Industry 4.0: Analysis of the digital transformation of energy efficiency programs," Applied Energy, 2025.
[11] D.Bousnina and G.Guerassimoff, "Optimal energy management in smart energy systems: A deep reinforcement learning approach and a digital twin case-study," Smart Energy, 2024.
[12] D.V.e.al., "Next-Level Energy Management in Manufacturing: Facility-Level Energy Digital Twin Framework Based on Machine Learning and Automated Data Collection," Energies, 2025.