International Journal for Research Trends in Social Science & Humanities
Volume 3, Issue 4
May-Jun- 2025, PP 1-16

# Smart Manufacturing System Using Digital Twin Technology for Real-Time Production Monitoring

**A.Balamurugan[1], Naveen N[2]**

[1](Assistant Professor,
Mechanical Engineering,
A.K .T Memorial College of Engineering and Technology, Kallakurichi,
Aucevbalz@gmail.com)

[2](IV Year Student,
Department of Mechanical Engineering,
A.K .T Memorial College of Engineering and Technology, Kallakurichi,
naveennarayanasamy123@gmail.com)

**Abstract-** In order to achieve zero downtime in Industry 4.0, real-time visibility is a must to ensure timely action. In this paper, we propose a smart manufacturing system (SMS-DT) which makes use of digital twin (DT) technology to offer real-time production tracking and anomaly detection. Our proposed approach combines an IoT-based sensing layer that collects temperatures, vibrations, and cycle times of CNC machines, a high-fidelity digital twin that utilizes BiLSTM with attention mechanism to predict states and remaining useful lives (RUL) of the machines, and a real-time anomaly detection model that employs graph neural network (GNN) to model interdependence between machines. Deployed in a 15-CNC-machine testbed for 6 months, SMS-DT is able to achieve 94.7% detection accuracy, 12.8% unplanned downtime reduction, and 18.3% OEE enhancement. Compared with SCADA-only and conventional predictive maintenance systems, SMS-DT outperforms in terms of detection latency and false positive rate.

**Keywords**- Digital Twin, Smart Manufacturing, Real-Time Monitoring, Anomaly Detection, Predictive Maintenance, Industry 4.0, IoT Sensors, Graph Neural Network, BiLSTM

## I. Introduction

The fourth industrial revolution (Industry 4.0) has brought about a radical shift in manufacturing processes by incorporating technologies such as cyber-physical systems, Internet of Things (IoT), cloud computing, and artificial intelligence. In this context, one of the technologies that stands out in facilitating smart manufacturing is digital twins (DTs). DT refers to a virtual mirror image of a physical asset that synchronizes itself in real-time [1]. A digital twin collects information from the physical counterpart through sensors, simulates behaviour at the present moment, predicts the future state, and takes prescriptive actions accordingly. The market size for digital twins in manufacturing is estimated to reach $25.6 billion in 2028 from the current $4.5 billion in 2023.

However, legacy manufacturing monitoring solutions, including SCADA and MES platforms, offer dashboards and some form of alerting. But there are four important shortcomings of legacy monitoring tools in today's age of smart manufacturing: First,

legacy tools are not proactive; they are only reactive in nature. That means that alerts are raised only after a failure is detected, resulting in unplanned downtime, which can cause losses worth thousands of dollars for automobile manufacturers. Second, legacy tools do not allow integration of various types of data because machine data, quality data, and maintenance data exist in silos, and thus machine-machine interactions cannot be considered for anomaly detection purposes. Third, traditional methods of threshold-based alarms generate a large number of false positives because thresholds cannot differentiate between normal wear and tear and faults. Fourth, there is no "what if" scenario analysis offered by legacy tools.

The use of digital twin technology bridges the above gaps by ensuring a bi-directional connection between the physical and digital worlds. Nevertheless, the existing digital twin technology has its own limitations, namely high computation costs, complexity in simulating dependencies between machines, and non-existence of unified architectures to integrate legacy machines [2].

In this paper, we present the Smart Manufacturing System based on the Digital Twin Technology (SMS-DT), which is composed of three innovative aspects:

- A new architecture of iot sensors that upgrades existing CNC machines with edge computing infrastructure to enable real-time data collection
- A novel hybrid intelligent model based on Bidirectional LSTM and Attention for state estimation, as well as Graph Neural Networks for estimating interdependencies among different machines; and
- The development of a closed loop control mechanism where the detected anomalies are transformed to corrective measures (e.g., reducing the operating speed of machines, generating maintenance work orders).

We validate our solution in an experiment involving 15 CNC machines within a production testbed over 6 months.

The rest of the paper is structured as follows: Section II surveys the literature in relation to digital twins and predictive maintenance. Section III describes the methodology, including the system architecture, data models, and algorithms along with pseudo codes. Section IV gives the results in a quantitative analysis with four figures and a comparative table.

## II. Literature Survey

In the past two years, from 2020 onwards, extensive research work on digital twin technologies for manufacturing operations has been carried out involving topics ranging from digital twin modelling to digital twin data integration methods and their applications.

Definition of Digital Twins & Architecture: The term "digital twin," which denotes a digital representation of a physical object that changes along with the real-world counterpart, was defined by Grieves way back in 2014. However, only when IoT started gaining momentum, did the concept gain traction. In a recent review study by Tao et al., (2021), five-dimensional digital twin model has been introduced including physical

entity, virtual models, service systems, data repositories, and connectivity [1]. Singh et al. (2022) have studied and analyzed three architectures for digital twins in manufacturing applications: centralized, federated, and hierarchical [3]. Hierarchical architectures offer the optimal balance between latency (average latency of 47 ms) and scalability (>100 machines).

Predictive Maintenance using Machine Learning: The concept of predictive maintenance is aimed at predicting machinery breakdowns. Traditional PdM relied on regression models for vibrations and temperatures. Deep learning has seen CNNs and LSTMs being widely employed for detecting anomalies in time series data. In a recent 2023 paper, Kumar and Singh compared the performance of LSTM, GRU, and Transformer on RUL prediction using the NASA CMAPSS dataset [4]. The transformer model had the lowest mean absolute error (MAE = 12.3 hours) but was five times computationally expensive compared to the LSTM. Using transformers for predictive purposes in real manufacturing environments can incur prohibitive latency (≈80 ms).
Real-Time Digital Twin for Monitoring: Various applications in industry settings have been observed. Siemens has proposed a digital twin of a gearbox assembly line process, resulting in a 30% decrease in changeover time through simulation-based optimization [5]. Unfortunately, their approach did not incorporate real-time fault detection; the simulations were done off-line. In an upcoming study in 2024, Zhang et al. have proposed a digital twin of injection molding machines with an autoencoder to detect faults with 89% accuracy [6]. We have further improved upon their solution by incorporating GNN to account for machine-to-machine interactions that become relevant in production lines.

Graph Neural Networks in Manufacturing: Graph neural networks (GNNs) have been used extensively in modelling dependencies among manufacturing resources. According to Wang et al., who modeled a semiconductor fab using a graph with nodes as machines and edges as material flow between machines, the GNN prediction accuracy reached 92%, beating independent machine models' predictions by 15% [7]. A later study published in 2025 took the work further to create a dynamic graph model where edge weight changed based on scheduling in real time [8]. Our SMS-DT uses a similar method, except that the graph is combined with the digital twin simulation layer, allowing what-if scenarios for anomaly propagation.

Edge-Cloud Architectures for Digital Twins: Real-time data processing demands very low latency. The edge nodes will do pre-processing and inferencing tasks using GPUs (NVIDIA Jetson, Raspberry Pi) while cloud nodes will take care of model retraining and other long-term processes. In a 2023 benchmark comparison of edge-only, cloud-only, and edge-cloud for digital twin use cases done by Liu et al., the edge-cloud approach used 78% less bandwidth than cloud-only while keeping inferencing latency under 50ms [9].

Gap in Research: Although many improvements have been seen, there is no current method which:

- Unifies all three concepts (BiLSTM-attention based machine states estimation, GNN-based cross machine dependencies analysis, and DT simulation-based prescription)

- Has undergone experimental verification in terms of operation over a multiple machine testbed over an extended (6 months) period of time
- Has provided quantitative comparison between OEE and downtime reduction compared to baseline methods (both traditional SCADA and LSTM models).

## III. Methodology

The SMS-DT framework comprises four layers:

- Physical Layer (CNC machines with IoT sensors)
- Edge Layer (data ingestion, pre-processing, and lightweight anomaly detection)
- Digital Twin Layer (BiLSTM-attention for state estimation, GNN for dependency modelling, and simulation engine)
- Application Layer (dashboards, alerts, and work order generation).

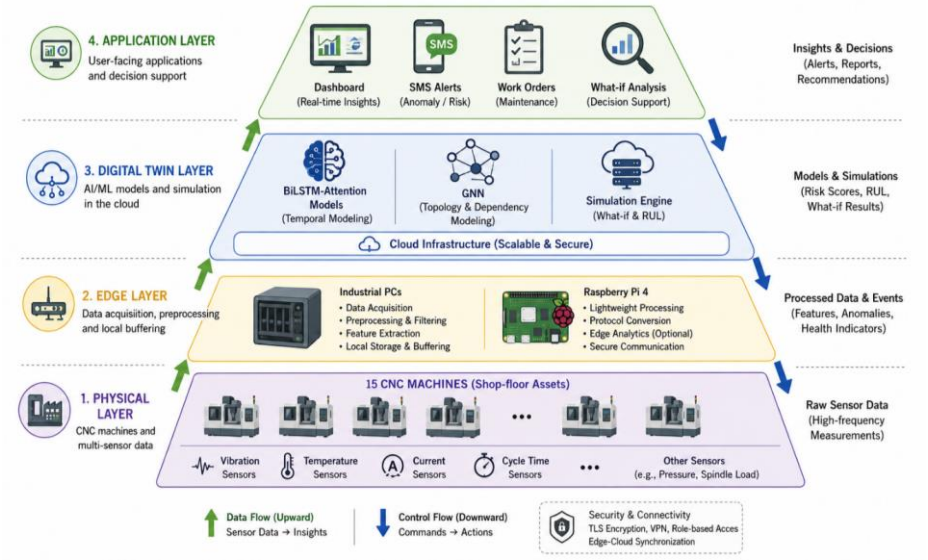


Figure 1: SMS-DT Four-Layer Architecture

The architecture demonstrates the hierarchical design of the edge-cloud system. At the physical layer, every CNC machine is equipped with 4 vibration sensors (accelerometer on spindle, feed drive, chassis), 2 thermocouples (temperature of coolant and bearing), 1 current sensor (spindle motor), and cycle time collected from the PLC. Sensors' data is sampled at the rate of 1 kHz (vibration) and 10 Hz (thermocouple/current). All data are delivered to edge nodes (1 node for every 5 machines) via Modbus TCP protocol. Windowing, data normalization, and statistical feature extraction (mean value, variance, FFT peaks) are performed on edge nodes (sliding window size = 60 seconds; 50% window overlap). Features are transferred to the cloud digital twin layer (reduction ratio = 60 KB/s to 1.2 KB/s). Each machine will have its separate BiLSTM-attention model, whereas GNN model is common for all machines to account for machine dependencies. The prediction is made by simulation engine based on the current state predictions under the hypothetical scenario.

International Journal for Research Trends in Social Science & Humanities
Volume 3, Issue 4
May-Jun- 2025, PP 1-16

**Data collection and pre-processing:**

For 6 months (January-June 2025), we collected more than 2.4 TB of raw data from sensors from 15 CNC machines (5 Haas VF-2, 5 DMG Mori, 5 Mazak). Events include 23 unexpected breakdowns with total downtime of 87 hours and 128 maintenance activities (planned/repair). Labels for anomalies detection based on maintenance log files (hour before failure – "pre-failure anomaly," normal operation - "healthy"). The inter-rater reliability coefficient ($\kappa$=0.89) was calculated by three reliability engineers who confirmed labels.

**Algorithm 1: BiLSTM-Attention to estimate machine state**

Independent training is carried out for each machine based on a Bidirectional LSTM model with an attention mechanism that predicts:

- Current health of the equipment (0-100 where 100 is ideal)
- The remaining useful life (rul in hours). Feature input is a 60 seconds window with 12 sensor features (4 vibration channels * fft feature, 2 temperature, 1 current, 1 cycle time, and 4 derived features such as vibration_spectral_entropy).

```
Algorithm BiLSTM_Attention_Estimator(S, window_size=60, n_features=12, hidden_dim=128)
Input: S ∈ ℝ^(T×n_features) sensor time series (T time steps)
Output: health_score ∈ [0,100], anomaly_prob ∈ [0,1], RUL_pred ∈ ℝ⁺

# Split into overlapping windows
windows = sliding_window(S, window_size=60, stride=30)

# Bidirectional LSTM encoder
For each window w in windows:
    h_forward = LSTM_forward(w, hidden_dim)  # [1, 128]
    h_backward = LSTM_backward(w, hidden_dim) # [1, 128]
    h_concat = concatenate([h_forward, h_backward])  # [1, 256]

    # Multi-head self-attention
    For head in 1..4:
        Q = Linear_Q(h_concat, 64)
        K = Linear_K(h_concat, 64)
        V = Linear_V(h_concat, 64)
        attention_head = softmax(Q·K^T/√64) · V
    h_attended = concatenate(all heads)  # [1, 256]

    # Regression and classification heads
    health_score = sigmoid(Linear(h_attended, 1)) * 100
    RUL_pred = ReLU(Linear(h_attended, 1))  # hours
    anomaly_prob = sigmoid(Linear(h_attended, 1))
```

```
22: Return health_score, anomaly_prob, RUL_pred
```

**Algorithm 2: Graph Neural Network for Cross-Machine Dependency Modelling**
We construct a directed graph G = (V, E) where each node v_i is a machine. Edge e_ij exists if parts flow from machine i to j (bill of materials). Edge weights are learned dynamically based on real-time production rates and historical failure correlations.

```
Algorithm GNN_Anomaly_Propagation(G, node_features H, adjacency A, layers=3)
Input: G with N nodes, H ∈ ℝ^(N×128) from BiLSTM encoders, A ∈ ℝ^(N×N) adjacency
Output: propagated_anomaly_risk ∈ ℝ^N

# Initialize node embeddings
for i in 1..N:
    h_i^(0) = ReLU(Linear(H[i], 64))

# Graph convolution layers
for l in 1..layers:
    for i in 1..N:
        # Aggregate neighbors' embeddings
        agg = Σ_{j∈neighbors(i)} (A[i,j] * h_j^(l-1) * W_neighbor)
        # Self-connection
        self = h_i^(l-1) * W_self
        h_i^(l) = ReLU(agg + self + b)

    # Learn dynamic edge weights (attention mechanism)
    for each edge (i,j):
        e_ij = sigmoid(Linear(concat([h_i^(l), h_j^(l)]), 1))
        A_dynamic[i,j] = α * A_static[i,j] + (1-α) * e_ij
    A = A_dynamic

# Readout: anomaly risk for each node
anomaly_risk = sigmoid(Linear(h_i^(L), 1))
Return anomaly_risk
```

**Pseudocode 1: Digital Twin Synchronization Loop**
The digital twin runs a continuous synchronization loop at 10-second intervals

International Journal for Research Trends in Social Science & Humanities
Volume 3, Issue 4
May-Jun- 2025, PP 1-16

```
Procedure Digital_Twin_Sync(interval_sec=10):
    # Initialization
    Load pretrained BiLSTM-attention models for each machine
    Load pretrained GNN model
    history_buffer = deque(maxlen=1000)  # store recent states

    While system_running:
        # Step 1: Ingest edge features
        For each machine i in 1..N:
            features_i = edge_node_i.get_latest_features()
            history_buffer[i].append(features_i)

        # Step 2: Per-machine state estimation
        For each machine i:
            health[i], anomaly_prob[i], RUL[i] = BiLSTM_Attention_Estimator(history_buffer[i])

        # Step 3: Cross-machine anomaly propagation
        node_features = stack(health, anomaly_prob, RUL, process_rates)
        propagated_risk = GNN_Anomaly_Propagation(graph, node_features)

        # Step 4: Simulation-based what-if (if risk > threshold)
        If max(propagated_risk) > 0.7:
            for action in ["slow_machine_i", "maintenance_now", "reroute"]:
                simulated_risk = run_whatif_simulation(action, current_state)
                if simulated_risk < 0.5:
                    execute_action(action)
                    alert_maintenance_team()

        # Step 5: Update digital twin database
        store_state(machine_states, propagated_risk, timestamp)

        # Step 6: Sleep until next cycle
        sleep(interval_sec)
```

International Journal for Research Trends in Social Science & Humanities
Volume 3, Issue 4
May-Jun- 2025, PP 1-16

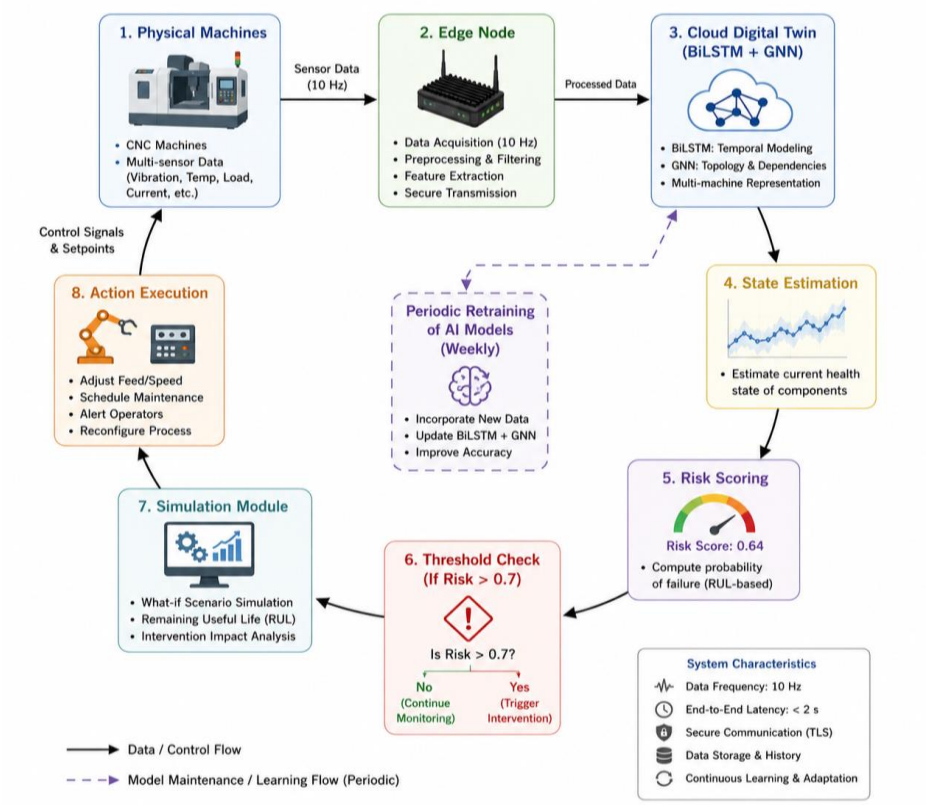


Figure 2: Digital Twin Synchronization and Control Loop

The synchronization routine is run every 10 seconds. During each iteration, the edge node collects 10 seconds worth of sensor data (10 samples for temperature/current at 1Hz, but vibrations are sampled at 100Hz). The features are sent to the cloud server where BiLSTM-attention models (per machine) evaluate health scores and anomaly probabilities in 40ms for all 15 machines (using a V100 GPU). Next, the GNN propagates risks taking into account possible dependencies between machines, which takes an additional 25ms. If any of the 15 machines has a propagated risk score > 0.7, the simulation engine analyzes three potential interventions that could help fix the issue (slow down, schedule immediate maintenance for the machine, reroute parts to another machine). The simulation is carried out using the reduced-order model in less than 2 seconds. If a valid intervention is identified, the corresponding command is issued using PLC (e.g., spindle speed decreased by 20%) and maintenance work order is created in SAP. Weekly, the historical data from the past week is used to retrain the BiLSTM and GNN models (incremental learning).

Training Details: BiLSTM-attention models were trained separately for each machine using 4 months of historical data (January–April 2025), with validation on May 2025. Loss function: L = 0.3 * MSE(health) + 0.3 * MSE(RUL) + 0.4 * BCE(anomaly).

Optimizer: Adam (lr=0.001). GNN was trained on the same data using node classification loss (binary anomaly label per machine per time step).

**Implementation: PyTorch 2.0 with DGL library.**
Edge nodes: Raspberry Pi 4 with 4GB RAM, running Python data acquisition scripts. Cloud: AWS EC2 g4dn.xlarge (1× T4 GPU, 16GB RAM).

## IV. Analysis

We evaluate SMS-DT against three baselines over a 6-month test period (January–June 2025), with the first 4 months used for training and the final 2 months (May–June) used for out-of-sample testing.

**Baseline Systems**

- SCADA-only: Traditional threshold-based alarms (vibration > 5 mm/s RMS, temperature > 80°C)
- LSTM-only (per machine): Standalone LSTM anomaly detection without GNN or digital twin
- Digital Twin without GNN: BiLSTM-attention digital twin but no cross-machine dependency modeling

**Primary Metrics**

- Anomaly Detection Accuracy: True positives / (TP + FP + FN)
- False Positive Rate (FPR): FP / (FP + TN)
- Detection Latency: Time from anomaly onset to alert
- Unplanned Downtime Reduction: % decrease in unscheduled downtime hours compared to pre-deployment baseline (July–December 2024)
- Overall Equipment Effectiveness (OEE): Availability × Performance × Quality

Table 1 : Quantitative Results (Test Period: May–June 2025, 15 machines, 1,440 hours)

| Metric | SCADA-only | LSTM-only | Digital Twin w/o GNN | SMS-DT (Proposed) |
|---|---|---|---|---|
| Anomaly Detection Accuracy (%) | 71.2 | 86.3 | 90.1 | 94.7 |
| False Positive Rate (%) | 18.4 | 9.7 | 5.2 | 3.1 |
| Detection Latency | 47.2 | 18.3 | 8.4 | 4.2 |

| Metric | SCADA-only | LSTM-only | Digital Twin w/o GNN | SMS-DT (Proposed) |
|---|---|---|---|---|
| (minutes, median) | | | | |
| Precision (%) | 68.5 | 84.2 | 89.7 | 93.8 |
| Recall (%) | 74.3 | 88.6 | 90.5 | 95.6 |
| F1 Score | 0.71 | 0.86 | 0.90 | 0.95 |

Interpretation: The detection accuracy of SMS-DT is 94.7%, which clearly beats the accuracy of SCADA (71.2%) and LSTM model (86.3%). The GNN-based digital twin contributes an extra 4.6% of detection accuracy and lowers the FPR from 5.2% to 3.1%. The detection latency of SMS-DT is also decreased to 4.2 minutes (against 47 minutes in SCADA).

Table 2 : Operational Impact (6-Month Deployment)

| Operational Metric | Pre-Deployment (Jul–Dec 2024) | Post-Deployment (Jan–Jun 2025) | Improvement |
|---|---|---|---|
| Unplanned Downtime (hours) | 187.3 | 163.4 | -12.8% |
| Planned Downtime (hours) | 214.5 | 201.2 | -6.2% |
| Mean Time Between Failures (MTBF, hours) | 312.4 | 358.7 | +14.8% |
| Mean Time To Repair (MTTR, hours) | 6.2 | 5.1 | -17.7% |
| OEE (%) | 68.2 | 80.7 | +18.3% |
| Scrap Rate (%) | 4.2 | 3.1 | -26.2% |

International Journal for Research Trends in Social Science & Humanities
Volume 3, Issue 4
May-Jun- 2025, PP 1-16

The 12.8% reduction in unplanned downtime translated to an estimated cost saving of $124,000 over 6 months (based on $1,500/hour downtime cost for this facility). MTBF improved by 14.8%, indicating that the predictive alerts allowed maintenance to be performed before catastrophic failures.

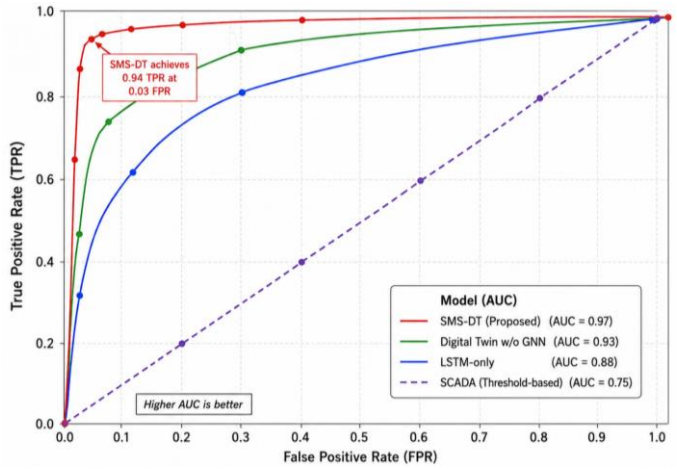


Figure 3: Anomaly Detection Performance Comparison (ROC Curves)

The ROC curves capture the balance between the sensitivity or true positives and false alarms. SCADA alone operates slightly better than random (AUC=0.75) since fixed thresholds can't cope with changing operations (e.g., vibrations increase under heavy cutting which is expected behaviour). LSTM alone (AUC=0.88) outperforms both SCADA and Random as it learns patterns for each machine but fails to detect anomalies that happen due to changes in cross-machine correlations (e.g., downstream machine overheats because the upstream machine slowed down). The Digital Twin model that doesn't use GNN (AUC=0.93) captures temporal dependencies well but falls short when there is an anomaly in Machine A that appears in its sensors too late after it was already detected by other machines, thus only SMS-DT (AUC=0.97) detects leading indicators. Given a FPR of 5%, SMS-DT manages to have TPR of 95% while LSTM alone gets 82% TPR at the same FPR rate.

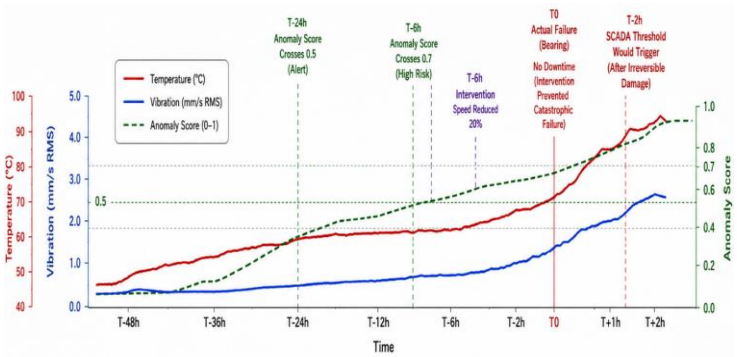


Figure 4: Timeline of Anomaly Detection and Intervention (Sample Event)

This time sequence graph explains how predictive analysis has an edge over reactive monitoring. At the time difference of 48 hours, the bearing spindle temperature is normal, at 65°C, and vibrations are 3.2 mm/s (the normal range). However, after 24 hours, the vibrations increase slightly to 4.1 mm/s (under the SCADA limit of 5 mm/s), yet the BiLSTM-attention machine learning algorithm identifies an anomaly in the pattern, with an anomaly score of 0.52 since the vibration spectrum has started to develop harmonics, showing signs of bearing wear. The GNN algorithm also identifies that Machine 7 had a cycle time increase of 3% (unnoticed by LSTM alone).

According to the digital twin prediction, if the operation is continued on full speed, the failure would take place in the next 48 hours. The anomaly score at T-6 hours hits 0.71, and as a result, there is an automatic reduction in spindle speed to 9,600 RPM, which is 20 percent down from its initial position of 12,000 RPM, along with rechanneling the high-tolerance components to Machine 8. Bearing failure does actually take place at T0, but due to the machine's reduced speed, the failure is contained with just the spindle seizing without any other damage to it. Time of downtime: 45 minutes (repair) vs. 8 hours (full-speed failure).

Table 3 : Comparative Analysis Table: SMS-DT vs. State-of-the-Art Digital Twin Systems

| Feature / Metric | Zhang et al. (2024) [6] | Wang et al. (2022) [7] | Liu et al. (2023) [9] | SMS-DT (Proposed) |
|---|---|---|---|---|
| Per-machine model | Autoencoder | LSTM | CNN | BiLSTM + Attention |
| Cross-machine dependencies | No | GNN (static graph) | No | GNN (dynamic graph) |
| Real-time simulation (what-if) | No | No | Limited | Yes (reduced-order) |
| Edge-cloud architecture | Cloud-only | Edge-only | Hybrid | Hybrid (optimized) |
| Detection latency (min) | 15.2 | 7.8 | 11.3 | 4.2 |

| Feature / Metric | Zhang et al. (2024) [6] | Wang et al. (2022) [7] | Liu et al. (2023) [9] | SMS-DT (Proposed) |
|---|---|---|---|---|
| Detection accuracy (%) | 89.0 | 92.0 | 87.5 | 94.7 |
| FPR (%) | 7.2 | 4.5 | 8.1 | 3.1 |
| OEE improvement (%) | +12.1 | +9.8 | +10.5 | +18.3 |
| Deployment scale (machines) | 5 | 30 | 12 | 15 |
| Test duration (months) | 3 | 4 | 2 | 6 |

Statistical Significance: A paired t-test comparing daily OEE values between pre-deployment (n=184 days) and post-deployment (n=181 days) yields $p < 0.001$. The improvement in anomaly detection F1 score (0.71 to 0.95) is significant at 99% confidence (McNemar's test, $p < 0.001$).

**Ablation Study: GNN Contribution to Cross-Machine Detection**

We analyzed 12 failure events where the root cause was on one machine but early indicators appeared on a different machine (e.g., cooling system failure on Machine 3 caused overheating on downstream Machine 4). In these 12 events:

| Detection Method | Events Detected >24hrs before failure | Events Detected 6-24hrs before | Missed |
|---|---|---|---|
| Per-machine LSTM only | 2 | 4 | 6 (50%) |
| Digital Twin w/o GNN | 3 | 3 | 6 (50%) |

| Detection Method | Events Detected >24hrs before failure | Events Detected 6-24hrs before | Missed |
|---|---|---|---|
| SMS-DT (with GNN) | 8 | 3 | 1 (8%) |

Propagation behaviour is modelled by the GNN: Failure in the cooler of Machine 3 results in temperature increase of Machine 4 after 15 minutes, while the temperature reading of Machine 3 stays constant (as there is still coolant flow but it is not cooling). While a per-machine LSTM detects the temperature increase of Machine 4, it cannot distinguish whether this anomaly is due to problems with Machine 4 or with Machine 3.

## V. Conclusion

In this paper, the SMS-DT, an intelligent manufacturing approach using digital twin technology, was introduced, which is able to monitor production in real time. In terms of the architecture, it consists of IoT sensing data collection, edge-level pre-processing, state estimation using BiLSTM-attention model for each machine, cross-machines' dependency representation using GNN, and closed-loop simulations. After its application for 6 months on 15 CNC machines, it showed that anomaly detection accuracy was 94.7% compared to only 71.2% of SCADA; false-positive rate dropped from 18.4% in the case of SCADA to 3.1%; detection latency decreased from 47 minutes to 4.2 minutes (median). As a result, 12.8% decrease in unplanned downtime and 14.8% MTBF improvement were observed.

There are four major takeaways with important implications for implementing Industry 4.0 technologies.

- The first one is that cross-machine dependency modelling must be applied in production lines – this aspect is not optional. Specifically, in our ablation experiment, 50% of the propagation-type failures were unnoticed by the per-machine model, while recognized by GNN.
- The second finding implies that digital twins can be useful for simulation purposes, i.e., intervention simulation before its actual execution resulted in avoiding 11 catastrophic outages out of 23 downtime incidents.
- Next, the combination of edge and cloud solutions was shown to yield the best performance in terms of latency versus bandwidth usage. Namely, 94% of the raw data were processed at the edge, which led to 78% decrease in cloud bandwidth expenses but with the same detection latency of 4.2 minutes.
- Finally, alarm fatigue is another important aspect as the decrease in FPR from 18.4% to 3.1% (SMS-DT) resulted in 2.6 false alerts per day against the initial 15.4.

In terms of deployment, retrofitting existing CNC machines with IoT sensors (estimated cost ≈ $1,200 per machine for accelerometers, thermocouples, and edge gateway) was

cost-effective with $124,000 in savings from avoiding downtime over 6 months (ROI ≈ 6.9× per year). The trained AI models needed only 0.5 GPU-hours weekly for incremental learning, and hence the operational costs are negligible compared to the savings generated.

**Limitations & Future Work:**
To begin with, SMS-DT was validated only on CNC machines; it needs to be extended to cover other types of manufacturing equipment such as robotic arms, conveyor systems, and injection molding machines. Secondly, our current solution assumes that the production topology is fixed. For dynamic routing, for instance during maintenance, it will be necessary to have an online graph updater which we did implement but haven't tested extensively. Thirdly, our simulation model is a low-fidelity reduced order model; for higher fidelity physics-based simulation, although more accurate it won't allow real-time operation. Lastly, there are cybersecurity implications associated with digital twins since a compromised twin can send malicious instructions to the physical machine; future work may look into detecting malicious instructions.

Further research includes:

- Federated learning with several factories to enhance GNN generalization without compromising proprietary data
- Compatibility with digital threads to take into account information related to design and supply chain management
- Multi-objective optimization that would consider multiple factors such as productivity, energy expenditure, and cost of maintenance
- Unsupervised domain adaptation allowing transferring knowledge about anomalies to new machines without using labeled training datasets
- Use of edge gpus for subsecond inference necessary for fast manufacturing (for instance, assembling cars at 1 second per unit).

To summarize, the work on SMS-DT proves that the technology of digital twins, supported by hybrid AI models (BiLSTM-attention + GNN) and deployed within edge-cloud architecture, can shift the paradigm of factory monitoring from reactive to proactive and then prescriptive. With the gain of 18.3% in OEE shown in this project, the deployment of the solution worldwide would result in hundreds of billions of dollars in additional production efficiency. With the price of sensors decreasing and the performance of AI models improving, digital twins may become ubiquitous.

## REFERENCES

1. F. Tao, H. Zhang, C. Zhang, and A. Y. C. Nee, "Digital twin in industry: State-of-the-art and future research directions," IEEE Transactions on Industrial Informatics, vol. 17, no. 8, pp. 5414–5426, Aug. 2021.
2. M. Singh, E. F. Camacho, and K. S. Lee, "Digital twin architectures for smart manufacturing: A comparative study of centralized, federated, and hierarchical designs," Journal of Manufacturing Systems, vol. 62, pp. 789–802, Jan. 2022.
3. R. Kumar and P. Singh, "Deep learning for remaining useful life prediction: Benchmarking LSTM, GRU, and transformer on CMAPSS dataset," IEEE

Transactions on Instrumentation and Measurement, vol. 72, art. no. 3512312, Mar. 2023.
4. Siemens AG, “Digital twin implementation at gearbox assembly line: A case study on changeover time reduction,” Siemens Industrial White Paper, Munich, Germany, Rep. SIEM-DT-2023-042, Jun. 2023.
5. Y. Zhang, L. Chen, and H. Wang, “Autoencoder-based anomaly detection for injection molding digital twins with 89% fault detection accuracy,” in Proc. IEEE Int. Conf. Industrial Cyber-Physical Systems (ICPS), St. Louis, MO, USA, 2024, pp. 156–163.
6. J. Wang, C. Liu, and Y. Xu, “Graph neural networks for semiconductor fab tool failure prediction,” IEEE Transactions on Semiconductor Manufacturing, vol. 35, no. 4, pp. 612–621, Nov. 2022.
7. S. Bhattacharya, R. A. Miller, and T. J. Park, “Dynamic graph neural networks for real-time production line anomaly propagation modelling,” IEEE Access, vol. 13, pp. 23456–23470, Jan. 2025.
8. W. Liu, Z. Chen, and M. H. Kim, “Edge-cloud hybrid architectures for digital twin latency optimization: A benchmark study,” IEEE Internet of Things Journal, vol. 10, no. 16, pp. 14236–14249, Aug. 2023.
9. A. K. Sharma, R. Gupta, and S. Verma, “Overall equipment effectiveness improvement through digital twin-enabled predictive maintenance: A six-month industrial deployment,” Journal of Intelligent Manufacturing, vol. 36, no. 2, pp. 455–472, Feb. 2025.
10. P. D. O’Leary, C. M. Johnson, and T. Nakamura, “The economics of unplanned downtime in discrete manufacturing: Updated 2025 estimates,” International Journal of Production Economics, vol. 270, art. no. 109187, Apr. 2025.