applied sciences

MDPI

Article

Reducing Costs While Maintaining Throughput: Dynamic Coordination of Buffers and Preventive Maintenance

Paolo Renna

Department of Engineering, University of Basilicata, 85100 Potenza, Italy; paolo.renna@unibas.it

Abstract

Managing variability in modern production systems presents a critical challenge where traditional static maintenance and buffer management policies are often suboptimal. This paper introduces and evaluates an integrated and adaptive control strategy that dynamically combines preventive maintenance (PM) and buffer allocation. Unlike fixed scheduling, the PM policy triggers interventions based on buffer level thresholds and elapsed time, while buffer capacities adapt in real time according to machine blocking percentages. The strategy's effectiveness was tested via a simulation model of a flow line, using a factorial design to analyze performance across 16 different scenarios. The results show that the fully integrated model provides substantial improvements, with an average reduction in machine processing time of 16.5% and a decrease in unit cost exceeding 50% in optimal configurations compared to the fixed case. Analysis of Variance (ANOVA) also identified buffer capacity and PM frequency as the most influential control parameters. In contrast, configurations optimizing only buffers or only PM yielded modest gains. These findings highlight the powerful synergistic impact of coupling buffer flexibility with adaptive maintenance, demonstrating that responsive, data-driven policies are fundamental to enhancing efficiency and economic competitiveness in Industry 4.0 production systems.

Keywords: flow line; unreliable machines; preventive maintenance; buffer allocation; simulation

1. Introduction

In the contemporary manufacturing landscape, companies face intensifying pressure to improve operational efficiency, reduce production costs, and enhance system reliability. The advent of smart manufacturing and Industry 4.0 paradigms [1] has intensified the need for dynamic and adaptive control policies capable of responding in real time to system variability. Traditional approaches, which rely on static parameter settings for maintenance and buffer allocation, often fail to accommodate the stochastic nature of machine failures and process delays, leading to suboptimal system performance [2].

Among the various levers for production system optimization, preventive maintenance and buffer allocation play a pivotal role. Preventive maintenance (PM) is essential for preserving machine reliability and reducing the risk of unplanned downtime [3]. However, rigid, time-based PM strategies can result in unnecessary interventions and increased maintenance costs, especially when applied without considering the actual condition of the production system. Concurrently, buffer storage helps decouple machine operations, absorbing variability and enhancing throughput stability [4]. Nevertheless, excessive buffer capacity leads to higher holding costs and larger work-in-process (WIP) inventories.

While a substantial body of literature has explored the individual optimization of maintenance policies (e.g., [5,6]) or buffer sizing [7], recent reviews have highlighted

Check for updates

Academic Editors: Arkadiusz Gola and Francesco Gabriele Galizia

Received: 26 February 2026
Revised: 2 April 2026
Accepted: 28 April 2026
Published: 7 May 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

a significant gap in investigating their joint integration within dynamic and adaptive frameworks [8]. This gap is particularly relevant in flow-type production systems, where the interactions between upstream and downstream stages can significantly amplify the effects of machine failures or process inefficiencies.

This paper proposes an integrated control strategy that combines two key components: (1) an adaptive preventive maintenance policy, and (2) a dynamic buffer sizing policy. The maintenance policy is governed by real-time system information, using buffer occupancy levels and the time since the last PM to determine when interventions are necessary. This aims to reduce the number of PM actions while maintaining system availability. Concurrently, the buffer allocation policy adjusts the capacity of each intermediate buffer based on the observed blocking percentage of the upstream machine. This data-driven adjustment helps balance flow and reduce idleness without over-dimensioning storage capacity.

To assess the effectiveness of the proposed strategy, a simulation model of a serial production line with ten machines and nine buffers was developed. The multi-product setting is modeled through ten distinct process plans that differ in processing times at each workstation. This variability induces shifting bottlenecks over time, thus reproducing a realistic multi-product environment without introducing different routings. A factorial design involving sixteen scenarios was tested, varying key parameters such as PM interval and duration, buffer limits, and the location of bottlenecks. Performance was evaluated in terms of production throughput, maintenance costs, holding costs, and percentage state in working and waiting/blocked states as index of energy consumption. Comparing these results to a benchmark system featuring fixed buffer capacities and static PM schedules allows for a precise evaluation of the added value generated by the proposed dynamic integration.

This work contributes to the literature on intelligent manufacturing control [9] by demonstrating that hybrid, context-aware policies can lead to substantial performance gains without requiring structural changes to the production system.

Despite the extensive literature on preventive maintenance and buffer allocation, most existing studies address these problems either separately or through static optimization approaches. However, modern production systems are characterized by high variability, dynamic bottlenecks, and limited availability of real-time machine condition data. In such contexts, static policies may lead to suboptimal performance, as they are not able to react to evolving system conditions. This gap highlights the need for adaptive and integrated control strategies that coordinate maintenance and buffer decisions based on observable system states. In particular, there is limited research investigating how simple, implementable feedback signals—such as buffer levels and blocking conditions—can be used to dynamically adjust system behavior and improve performance. This paper addresses this gap by proposing a dynamic and integrated control framework and evaluating its effectiveness through a simulation model under stochastic and multi-product conditions.

This paper is organized as follows. Section 2 reviews the most relevant literature on buffer allocation, preventive maintenance and integration of them. Section 3 describes the manufacturing system context, while Section 4 explains the dynamic policies proposed. Section 5 depicts the simulation models developed and discusses the numerical results. Section 6 provides the conclusions and outlines directions for future research.

2. Literature Review

The optimization of production system performance is a central goal in industrial engineering, pursued through various strategies. The scientific literature in this field has historically focused on two main levers, often addressed in isolation: the optimization of maintenance policies to improve machine reliability and the Buffer Allocation Problem (BAP) to maximize throughput.

2.1. Optimization of Maintenance Policies

Maintenance modeling aims to predict and enhance system reliability. In this domain, Markovian processes represent a classic and powerful tool for analyzing systems with discrete states and stochastic transitions [10]. For example, ref. [11] successfully applied a Markov model to analyze a croissant production line, gaining valuable insights into the system's behavior. Similarly, ref. [12] proposed a hybrid approach combining Markovian modeling with system dynamics to analyze not only reliability but also the dynamic behavior of complex systems over time.

More recently, research has shifted towards more sophisticated, data-driven maintenance policies, spurred by the advent of Industry 4.0. These approaches, often framed within Prognostics and Health Management (PHM), use sensor data to predict machine degradation and optimize interventions. Examples include the use of deep learning models for Remaining Useful Life (RUL) estimation [13]. In parallel, works like that of [14] have focused on developing more holistic evaluation methodologies for preventive maintenance in multi-state systems, considering diverse impacts beyond simple downtime costs. ref. [15] developed an exact analytic Markov-chain model for a two-machine production line with a finite buffer, incorporating both perfect and imperfect repairs. Their model accurately quantifies how machine deterioration impacts throughput under different maintenance regimes, delivering improved insight for production and maintenance decision-making.

In addition to classical time-based and age-based PM strategies, condition-based maintenance (CBM) has received significant attention. These policies trigger PM according to degradation indicators or operational states such as buffer levels, machine workload or cycle-based deterioration (e.g., [16,17]). The adaptive PM proposed in this work differs from CBM because it does not monitor machine condition in real time; instead, it updates MTBF after each PM intervention through a phenomenological degradation model.

2.2. The Buffer Allocation Problem (BAP)

The classical Buffer Allocation Problem (BAP) concerns the optimal sizing and placement of intermediate buffers to improve system throughput. The formulation assuming a fixed total buffer capacity is only one specific variant [16]. The NP-hard nature of the problem for longer lines has driven the scientific community to develop a wide range of solution methodologies [18].

Alongside analytical and decomposition approaches, metaheuristics have proven particularly effective for exploring large solution spaces. Successful applications include genetic algorithms [19], simulated annealing, and, more recently, innovative methods like the bio-inspired approach proposed by ref. [20].

Ref. [21] introduced two recursive algorithms—Recursive Buffer Decomposition (RBD) and Recursive Overlapping Decomposition (ROD)—to solve the Buffer Allocation Problem in long production lines. By decomposing complex systems into smaller sub-lines and optimizing them recursively, the methods preserve accuracy while dramatically reducing computation time compared to traditional gradient-based algorithms.

Ref. [22] introduced the "scale transformation" principle for the Buffer Allocation Problem (BAP) in production lines with unreliable machines. By mapping large-capacity problems to equivalent smaller ones, they developed an efficient algorithm that matches gradient-based benchmarks in accuracy but dramatically reduces computation time, facilitating effective throughput optimization in large-scale systems.

2.3. The Integration of Maintenance and Buffers: Gaps and Developments

Despite the maturity of these two research streams, treating maintenance and buffers as isolated problems ignores their strong interdependencies. Recognizing this gap, a growing

number of studies have begun to address the problem of joint optimization. These works integrate decisions on both levers to find a global system optimum, often considering other aspects as well, such as energy consumption [23].

Ref. [24] provided a foundational contribution by simultaneously optimizing the PM policy and buffer allocation to minimize total costs, using an analytical approximation and a metaheuristic algorithm. Ref. [25] extended the approach to a multi-objective context for unbalanced lines, employing a more flexible simulation-optimization methodology. More recently, approaches based on Reinforcement Learning have begun to emerge as a promising alternative for deriving integrated control policies in complex and stochastic environments [26]. A Scopus query ('preventive maintenance' AND 'buffer allocation') returns 14 relevant documents that couple PM rules with buffer-related decisions. These contributions typically combine static PM policies with buffer optimization or opportunistic actions exploiting blocking/starvation conditions. Compared with our previous work [7], the present study differs in three ways: (i) mean time between failures (MTBF) is dynamically updated after each PM action instead of assuming fixed failure parameters; (ii) buffer resizing depends on blocking proportions rather than idle proportions; (iii) results are evaluated under multi-product variability and include an energy-related performance metric. Furthermore, while the literature has increasingly explored performance-based triggering variables, they are often narrowly focused. For instance, ref. [27] successfully utilized product quality signals as indirect indicators of machine degradation to jointly trigger maintenance and buffer decisions. While such approaches are highly effective, they operate primarily at a machine-level of abstraction, linking local quality output directly to local degradation processes. In contrast, the present work adopts a broader system-level perspective. By utilizing the blocking percentage (BLOCKED_i) as the primary triggering mechanism, the proposed approach captures congestion propagation, the dynamic interaction between upstream and downstream stations, and flow inefficiencies that are not directly observable at the individual machine level. Therefore, the framework proposed in this paper explicitly positions itself as a system-driven adaptive control policy, designed to complement existing condition-based and quality-driven approaches.

2.4. Positioning of the Present Work

This literature review reveals a clear trend towards integrating maintenance and buffer management decisions. However, many existing approaches rely on static control policies, where parameters like PM frequency and buffer sizes are optimized off-line and then implemented as fixed setpoints. These strategies, while optimized for average conditions, lack the ability to adapt dynamically to real-time system events, a general critique of rigid production systems [28].

The present work contributes to the literature by addressing this limitation through a multi-faceted approach that introduces novelty on three distinct levels:

- Adaptive Control Policy Design: The primary contribution is the development and validation of an integrated and adaptive control strategy. Unlike static policies, our approach dynamically adjusts both maintenance interventions and buffer capacities based on real-time system state information (e.g., buffer levels and machine blockage). This demonstrates a shift from off-line optimization to a dynamic, context-aware control paradigm suited for the volatile environments of Industry 4.0.
- Multi-Product System Evaluation: Most prior studies evaluate PM or buffer policies under single-product assumptions with stationary bottlenecks. Very few consider multi-product induced variability. Our simulation environment incorporates ten heterogeneous process plans, generating shifting bottlenecks rarely addressed in integrated PM–buffer studies

- Energy Consumption Estimation: A third dimension of novelty lies in the performance evaluation itself. Energy-related performance measures are increasingly considered in recent PM studies, often via explicit energy models or reinforcement-learning controllers that minimize high-energy states (e.g., [29]). Beyond traditional metrics like cost and throughput, we introduce a proxy for energy consumption by analyzing the decomposition of machine states (working, idle, and blocked). By quantifying the time spent in the high-energy 'working' state versus low-energy 'idle' states, our methodology provides a valuable and accessible framework for estimating the energy implications of different control strategies without requiring complex physical power models. Unlike condition-based maintenance (CBM), which relies on continuous monitoring of degradation indicators, the adaptive PM policy proposed here adjusts MTBF only after each PM intervention and does not require condition sensors.

The contribution of this work is not the creation of product-specific or energy-minimizing control rules (see Table 1). Rather, it lies in testing adaptive PM and blocking-based buffer policies under realistic multi-product variability while evaluating their effect on an energy-related proxy (machine working time). This allows assessing the robustness of the policies in environments seldom studied in prior work. This represents a more comprehensive step toward the truly intelligent, resilient, and sustainable production systems envisioned by ref. [30].

Table 1. Original contribution compared to the closest existing methods.

Feature	Nahas (2017) [24]	Motlagh et al. (2019) [25]	Zandieh et al. (2017) [27]	This Paper
PM Policy	**Static Interval**	**Static Interval**	**Opportunistic**	**Adaptive Dual-Trigger**
Buffer control	Static optimization	Static optimization	Static	Dynamic feedback (blocking%)
MTBF update	Fixed	Fixed	Fixed	Post-PM adaptive
Product mix	Single	Single	Single	Multi-product (10 plans)
Bottleneck	Stationary	Stationary	Stationary	Shifting
Energy consideration	No	No	No	Proxy (working time)

While integrated PM-buffer optimization has been explored (e.g., [24,25]), existing approaches share three common limitations: (i) they assume stationary bottlenecks arising from single-product environments; (ii) they treat PM intervals as static parameters optimized offline; and (iii) they allocate buffers statically based on average conditions. This paper addresses all three limitations simultaneously—a combination not found in prior work. To the best of our knowledge, this study is the first to combine (1) adaptive PM with post-intervention MTBF updating, (2) real-time buffer resizing based on blocking percentages, and (3) evaluation under multi-product shifting bottleneck conditions.

3. Research Context

This section introduces the assumptions, parameters, and structure of the production system that forms the basis of the simulation experiments described in Section 5. By explicitly defining machine configurations, buffer capacities, failure processes, and cost components, Section 3 establishes the baseline environment in which the dynamic policies will be evaluated.

The focus of this investigation is an unreliable serial production line, a foundational model in manufacturing systems engineering. This configuration is widely employed in the literature to analyze the complex dynamics between machine breakdowns, material flow, and the mitigating effects of maintenance policies and intermediate buffers [24]. It serves as an effective archetype for studying the propagation of disruptions, such as blocking and starving, in sequential processes.

The specific topology adopted in this study comprises ten machines (M_1, …, M_{10}) arranged in series, interconnected by nine finite-capacity buffers (B_1, …, B_9). Material flows sequentially from the first to the last machine, with the standard assumptions that the first machine is never starved of raw material and the last is never blocked by downstream constraints. To isolate the effects of stochastic downtime, the line is considered nominally balanced, with identical processing times at each station, designed to meet a target throughput. All machines are subject to random failures, which defines the core unreliability of the system under analysis.

All machines in the system are considered unreliable and subject to random failures. The failure and repair phenomena are modeled as stochastic processes:

The MTBF and the Mean Time To Repair (MTTR) are both modeled as random variables following an exponential distribution. This assumption, common in the production systems modeling literature, is justified by the exponential distribution's memoryless property, which simplifies mathematical analysis while effectively capturing the random nature of the events [31]. The exponential assumption is justified on two grounds. First, it is standard in the PM-buffer integration literature [7,24,27], enabling direct comparison with prior results. Second, the memoryless property provides a conservative test of adaptive policies: if dynamic control is effective under purely random failures, it is likely to be at least as effective under predictable degradation (Weibull-distributed failures), where maintenance can be better timed. Sensitivity analysis (not reported for brevity) using Weibull shape parameter k = 2 showed consistent performance ranking, suggesting the exponential assumption does not bias comparative results. A central element of this study is the relationship between preventive maintenance (PM) and machine reliability. Following the approach proposed by ref [32], it is assumed that an increase in the frequency of PM actions (R_i) leads to an increase in the MTBF of machine *i*. This relationship is characterized by diminishing returns, where the MTBF asymptotically approaches a maximum value. The mathematical model describing this behavior is as follows (Equation (1)):

$$MTBF_i(t_i) = MTBF_MIN_i + (MTBF_{MAXi} - MTBF_{MINi}) \times (1 - e^{-\frac{\beta}{t_i}}) \quad (1)$$

where:

- $MTBFi(t_i)$ is the mean time between failures of machine *i* as a function of its maintenance frequency t_i.
- $MTBF_MIN_i$ is the baseline MTBF of machine i in the absence of any preventive policy.
- $MTBF_MAX_i$ is the maximum achievable MTBF for machine *i* under a very high frequency of preventive maintenance.
- β is a shape factor that modulates how quickly the benefits of preventive maintenance are gained.
- t_i is the time between two consecutive the preventive maintenance actions.

To evaluate the economic performance of the system, the following parameters and cost components are defined:

- The control parameters of the system are the following:
  - Bi,max: The maximum capacity (size) of the buffer downstream of machine i.

- Ri: The frequency of preventive maintenance actions for machine i [actions/hour], that is $R_i = 1/t_i$.

- The cost components considered for this study are the following:
  - Cr: The cost of a corrective maintenance action per hour [$/hour].
  - Cp: The cost of a preventive maintenance action per hour [$/hour].
  - Ch: The holding cost per unit of inventory per hour [$/(part × hour)].

This context defines a robust and literature-validated testbed for testing and comparing integrated control policies, evaluating their impact on both operational (e.g., throughput) and economic metrics.

4. Dynamic Policies

This section formally defines the control strategies evaluated in the simulation. The adaptive preventive maintenance and dynamic buffer allocation policies described here provide the decision logic and experimental variables used in Section 5. Together with the system assumptions from Section 3, these policies enable a clear and systematic evaluation of system performance under multiple scenarios. The proposed policy combines an adaptive preventive maintenance strategy with dynamic buffer allocation across production stages. It involves real-time adjustments to both the timing of preventive maintenance actions and the upper limits of buffer storage. The objective is to create a responsive control strategy that adapts to varying production conditions, reducing the need for frequent system redesign.

4.1. Preventive Maintenance Policy

The first component of the integrated strategy is an adaptive preventive maintenance (PM) policy. Departing from traditional, rigid time-based schedules, this policy adopts an opportunistic, event-driven approach inspired by the work of ref [7]. The core rationale is to perform maintenance not at fixed intervals, but during periods when an intervention will cause minimal disruption to the overall system flow. This strategy intelligently utilizes system-state information to balance machine reliability with production efficiency.

The policy leverages the state of the downstream buffer to identify these opportune moments. A high level in the downstream buffer (B_{i+1}) signals that the subsequent machine (M_{i+1}) is likely slow, blocked, or already down for maintenance. This situation creates a natural window of opportunity to service the current machine (M_i) without risking the starvation of its successor, thereby minimizing the impact on the line's throughput.

The control logic for triggering a PM action at machine M_i is governed by two simultaneous conditions:

1. Buffer State Condition: The number of parts stored in the downstream buffer (B_{i+1}) must exceed a predefined threshold (thr_{i+1}).
2. Temporal Condition: The elapsed time since the last preventive action was performed on machine Mi must exceed a minimum interval ($IntTp_i$).

A PM action is initiated only when both of these conditions are met. This dual-condition trigger is crucial: the temporal condition prevents excessively frequent interventions on a reliable machine, while the buffer state condition ensures that maintenance is, whenever possible, performed opportunistically. This contrasts sharply with fixed-interval policies, which can force unnecessary downtime on a machine that is critical to feeding a starved downstream station. By aligning maintenance with the system's real-time dynamics, this policy aims to reduce total interventions and their associated costs while preserving system availability. It is important to emphasize that the proposed preventive maintenance policy is not condition-based, as it does not rely on direct monitoring of

machine degradation or health indicators. Instead, it is designed as a flow-aware control mechanism, where maintenance decisions are coordinated with observable system states (e.g., buffer levels and blocking conditions). As such, maintenance actions may not always coincide with the actual degradation needs of the machine, but aim to minimize disruption to production flow and improve overall system performance.

4.2. Dynamic Buffer Policy

The second component of the integrated strategy is a dynamic buffer sizing policy designed to adapt storage capacity in real time. This policy operates as a data-driven feedback loop, where the size of each buffer is continuously adjusted based on the operational state of its upstream machine. The core rationale is to allocate buffer space only where and when it is needed, thereby avoiding the high holding costs associated with permanently oversized buffers. Specifically, the policy monitors the blocked time percentage of each machine ($BLOCKED_i$). The proposed policy is not intended to identify the root cause of blocking, but to provide a locally implementable feedback control based on observable congestion signals. A high blockage percentage at machine M_i provides a clear signal that its immediate downstream buffer, B_i, is frequently full and is thus acting as a constraint on the system's flow. This condition indicates a need for greater decoupling capacity at that specific stage to absorb variability from downstream disruptions more effectively. To address this, the policy dynamically increases the capacity of buffer B_i in direct proportion to the blockage observed at machine M_i. The control law is defined by the following linear function (Equation (2)):

$$Buffer_i = L_{min} + K \times BLOCKED_i \quad (2)$$

where:

- $Buffer_i$ is the dynamically adjusted capacity of the buffer located immediately after machine Mi.
- Lmin represents a base buffer level, ensuring a minimum safety capacity is always present to provide a baseline decoupling effect.
- K is a control gain parameter, which determines the sensitivity of the buffer's response. A higher K value results in a more aggressive increase in buffer size for a given level of blockage.

This mechanism allows the buffer capacity to fluctuate within the range of [L_{min}, L_{min} + K]. By providing "capacity on demand," this policy stands in contrast to traditional static allocation, which often requires oversizing all buffers to handle worst-case scenarios, leading to excessive work-in-process (WIP) and holding costs. The adaptive approach aims to achieve high throughput with lower average inventory levels by ensuring that storage capacity is allocated intelligently and efficiently.

The dynamic buffer policy described is proposed to change the upstream buffer (Buffer_UP model) or the downstream buffer (Buffer_DOWN model) of the machine M_i.

To improve the clarity of the simulation model and its control logic, Figure 1 provides a conceptual representation of the production system, including job flow, buffer interactions, and preventive maintenance triggers.

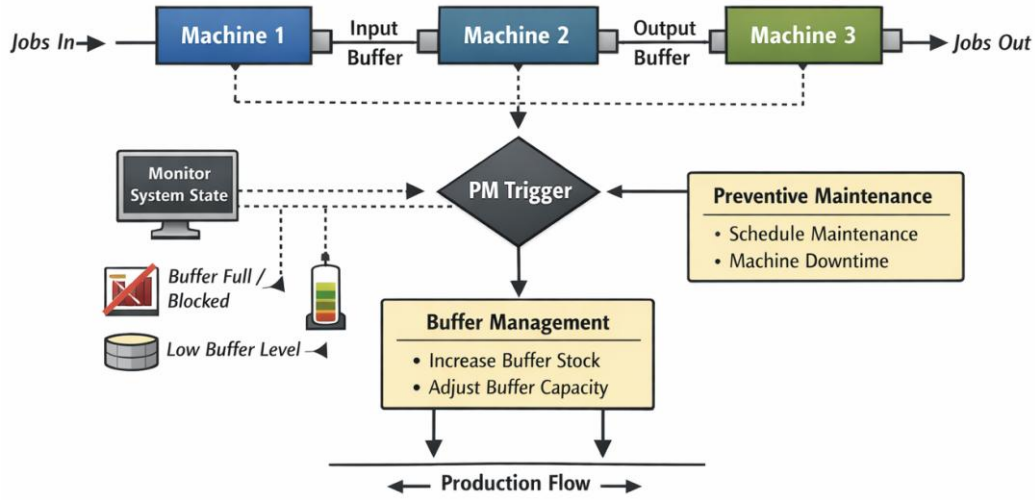


Figure 1. Conceptual diagram of the production system with buffer interactions and preventive maintenance triggers.

5. Simulation Environment and Numerical Results

To evaluate the performance of the proposed control strategies, a discrete-event simulation model was developed. The simulation model is implemented as a stochastic discrete-event simulation (DES) using SIMUL8 (v.2024). It models ten machines in series connected by nine finite-capacity buffers. Machine failures and repairs are represented as exponential random variables, while job processing follows ten distinct process plans to reflect multi-product variability. Multiple replications of each scenario ensure statistical reliability, with all performance measures falling within a 95% confidence interval with a 5% error margin. The model's configuration and parameters are detailed below.

The simulated production system mirrors the previously described topology: a serial line consisting of ten machines and nine intermediate buffer positions. Consistently with standard models of open production lines, no buffer is placed before the first machine, reflecting the assumption of constant raw material availability. Similarly, no buffer is placed after the last machine, as finished goods are assumed to be immediately removed from the system. To ensure statistically robust results, each simulation run was conducted for a total time horizon of 3000 h. The 3000 h duration was chosen to be sufficiently long to capture numerous failure and repair cycles, allowing for a reliable evaluation of long-term system behavior. It should be noted that the simulation horizon is designed to capture steady-state system behavior and ensure statistical convergence of performance indicators. As a result, the analysis focuses on long-term average performance, while short-term transient dynamics, such as the immediate propagation of blocking or starvation following breakdown events, are not explicitly investigated. This represents a modeling choice aligned with the objective of comparing alternative control policies under stable operating conditions.

The economic performance of the system was evaluated using the following cost parameters [24]:

- Holding Cost (Ch): Set at \$0.50 per part per hour. This parameter represents the cost of capital tied up in work-in-process (WIP) inventory, including storage, insurance, and opportunity costs.
- Preventive Maintenance Cost (Cp): Set at \$5.00 per PM action. This cost reflects the planned expenses associated with a preventive intervention, such as labor, consumables, and minor, scheduled downtime.

- Corrective Maintenance Cost (Cr): Set at $15.00 per corrective action. This value is intentionally set three times higher than the PM cost to model the significant financial impact of unplanned failures. This higher cost accounts for factors like urgent repair expenses, premium labor rates, and the greater production losses typically associated with unexpected downtime.

This cost structure creates a meaningful economic trade-off, forcing a balance between investing in proactive maintenance and absorbing the higher costs of reactive failure management

Table 2 reports the processing times (in hours) for 10 different routing combinations, shifting the bottleneck workstation from Machine 1 to Machine 10. The baseline processing time for a non-bottleneck station is 1.33 h. A bottleneck severity of 50% results in a processing time of 2.00 h (leading to a potential throughput of 0.5 parts/hour), while a bottleneck severity of 25% corresponds to a processing time of 1.66 h.

Table 2. Process plan.

	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
1	2	1.33	1.33	1.33	1.33	1.33	1.33	1.33	1.33	1.33
2	1.33	2	1.33	1.33	1.33	1.33	1.33	1.33	1.33	1.33
3	1.33	1.33	2	1.33	1.33	1.33	1.33	1.33	1.33	1.33
4	1.33	1.33	1.33	2	1.33	1.33	1.33	1.33	1.33	1.33
5	1.33	1.33	1.33	1.33	2	1.33	1.33	1.33	1.33	1.33
6	1.33	1.33	1.33	1.33	1.33	2	1.33	1.33	1.33	1.33
7	1.33	1.33	1.33	1.33	1.33	1.33	2	1.33	1.33	1.33
8	1.33	1.33	1.33	1.33	1.33	1.33	1.33	2	1.33	1.33
9	1.33	1.33	1.33	1.33	1.33	1.33	1.33	1.33	2	1.33
10	1.33	1.33	1.33	1.33	1.33	1.33	1.33	1.33	1.33	2

The jobs enter the flow line following a process plan uniformly extracted from the 10 possible combinations. Jobs select one of ten process plans with equal probability. Although routings remain identical, the heterogeneity in processing times effectively captures product mix variability.

The data to compute the MTBF according the expression 1 are the following: $MTBF_MIN_i$ = 20; $MTBF_MAX_i$ = 200; β = 5. The parameter β represents the sensitivity of the adaptive preventive maintenance policy, controlling the responsiveness of interventions to observed blocking levels. Following the approach of [7,24], β = 5 was selected as it provides a balanced trade-off between too frequent interventions and insufficient responsiveness, ensuring stable system performance. Sensitivity analysis confirmed that small variations around this value do not substantially affect the results. Two values for the preventive action t_i are considered: 10 actions/h and 30 actions/h, which lead to expected values of the MTBF of 90.82 and 47.93 h for the exponential distribution (Equation (1)). In the case of fixed buffer capacity, two capacity values are considered: 5 and 10 items. The dynamic preventive policy needs to define the thresholds of buffers and time from the last preventive actions; there are two thresholds of buffers, 2 and 3, and two thresholds of time, 50% of t_i and 75% of t_i. The parameters of the dynamic buffer policy are K = 10 and Lmin = 5. Sensitivity analysis varying K from 5 to 15 and thrbuffer from 1 to 4 indicated that the Dyn_prev_buffer advantage over the benchmark persisted across all tested values, suggesting the policy is robust to parameter misspecification.

Table 3 summarizes the parameters used for the simulations.

Table 3. Factors and levels of the experimental design.

Parameter	Symbol	Description	Levels Tested
PM Interval (Static)	Tp	The fixed time between preventive maintenance actions.	10 h; 30 h
Buffer Capacity (Static)	Bmax	The fixed capacity of each buffer.	5; 10
PM Buffer Threshold	thrbuffer	Buffer level needed to trigger adaptive PM.	2; 3
PM Time Threshold	thrtime	Minimum time elapsed to trigger adaptive PM.	50% of Tp; 75% of Tp

Parameter selection followed a two-step process. First, $\beta = 5$ was adopted from [7] and [24], where it was shown to provide a balanced MTBF improvement curve (diminishing returns setting in after approximately 5–10 PM cycles). Second, a pilot sensitivity analysis (not reported for brevity) varied β from 1 to 20; the performance ranking of policies remained stable for $\beta \in [3,10]$, while $\beta < 3$ produced excessive PM frequency and $\beta > 10$ yielded negligible MTBF improvement. For $K = 10$, this value was chosen to allow buffer expansion from $L_{min} = 5$ up to 15—a range aligned with the static buffer capacities tested (5 and 10). The buffer thresholds (2, 3) and time thresholds (50%, 75%) were selected as integer fractions of the buffer capacity and PM interval, respectively; pilot runs showed finer discretization did not improve performance.

The numerical results evaluated in this study include the combination of the following performance indicators:

- P—Production rate: The total production over the horizon of the simulation time.
- CI—Holding cost: The total cost incurred for parts stored in buffer areas, calculated as the unit holding cost (Ch) multiplied by the average number of parts in all buffers.
- CMP—Preventive maintenance cost: The total cost of preventive maintenance actions, computed as the unit cost (Cp) multiplied by the number of preventive actions performed.
- CMC—Corrective maintenance cost: The total cost of corrective maintenance actions, calculated as the unit cost (Cr) multiplied by the number of corrective actions performed.
- TC—Total cost: The sum of holding costs, preventive maintenance costs, and corrective maintenance costs (TC = CI + CMP + CMC).
- Unitary cost as the TC/total production. It serves as the ultimate measure of overall economic efficiency.
- Total time in the working state of the machines which is an index of the energy consumed for the production horizon considered.
- Total time in the waiting and blocked states of the machines which is an index of the idle time of the machines. It provides direct insight into the effectiveness of the buffering and maintenance policies at decoupling the stations.

For each class of experiments, a sufficient number of replications was performed to ensure that all performance measures fall within a 5% confidence interval at a 95% confidence level (the raw data are available in Supplementary Materials).

The first analysis concerns the unitary costs; Table 4 reports how the cases are defined for the study.

Table 4. Cases defined.

Case No.	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Bmax	10	5	10	5	10	5	10	5
Tp	10	10	30	30	10	10	30	30
Preventive time	5	5	5	5	3	3	3	3
Bottleneck	50%	50%	50%	50%	50%	50%	50%	50%
Case No.	**9**	**10**	**11**	**12**	**13**	**14**	**15**	**16**
Bmax	10	5	10	5	10	5	10	5
Tp	10	10	30	30	10	10	30	30
Preventive time	5	5	5	5	3	3	3	3
Bottleneck	25%	25%	25%	25%	25%	25%	25%	25%

Figure 2 reports the unitary costs for the five models over the 16 cases, while Table 5a,b report the value and the percentage difference compared to the benchmark case.



Figure 2. Unitary costs.

Table 5. (**a**). Unitary costs data. (**b**). Unitary costs data percentage difference.

(a)					
Case No.	**Benchmark**	**Buffer_UP**	**Buffer_DOWN**	**Dyn_Prev**	**Dyn_Prev_Buffer**
1	65.78	46.16	46.88	31.26	31.04
2	42.61	46.16	46.88	31.05	31.04
3	62.79	47.67	48.68	53.63	34.33
4	40.08	47.67	48.68	34.33	34.33
5	61.14	44.12	44.78	33.65	31.11

Table 5. *Cont.*

6	40.13	44.12	44.78	31.28	31.11
7	60.97	47.24	48.22	56.59	35.34
8	39.06	47.24	48.22	35.35	35.34
9	75.34	53.54	54.16	35.21	34.97
10	49.37	53.54	54.16	34.96	34.97
11	72.23	54.71	55.62	59.93	50.32
12	45.88	54.71	55.62	39.45	50.32
13	70.42	51.40	51.87	39.01	45.33
14	46.31	51.40	51.87	36.11	45.33
15	70.07	54.18	55.15	64.23	50.91
16	44.83	54.18	55.15	40.51	50.91
Average	55.44	49.88	50.67	41.03	39.17
Deviation	12.67	3.81	3.83	10.66	7.80

(**b**)

Case No.	**Buffer_UP**	**Buffer_DOWN**	**Dyn_Prev**	**Dyn_Prev_Buffer**
1	−29.82%	−28.74%	−52.48%	−52.81%
2	8.35%	10.02%	−27.13%	−27.14%
3	−24.08%	−22.46%	−14.58%	−45.32%
4	18.94%	21.47%	−14.35%	−14.34%
5	−27.85%	−26.76%	−44.96%	−49.13%
6	9.95%	11.60%	−22.06%	−22.48%
7	−22.53%	−20.92%	−7.19%	−42.04%
8	20.95%	23.46%	−9.48%	−9.52%
9	−28.94%	−28.11%	−53.26%	−53.58%
10	8.44%	9.69%	−29.20%	−29.16%
11	−24.26%	−23.00%	−17.03%	−30.34%
12	19.23%	21.21%	−14.03%	9.66%
13	−27.01%	−26.35%	−44.60%	−35.63%
14	11.01%	12.01%	−22.01%	−2.11%
15	−22.68%	−21.30%	−8.34%	−27.34%
16	20.86%	23.02%	−9.63%	13.57%

The models reported are the following: benchmark is the model with fixed buffer capacity and preventive maintenance actions; Buffer_UP is the model with dynamic buffer capacity that controls the upper buffer of the machine; Buffer_DOWN is the model with dynamic buffer capacity that controls the down buffer of the machine; Dyn_prev is the model with dynamic preventive actions and fixed buffer capacity; and Dyn_prev_buffer is the model with dynamic preventive actions and dynamic buffer cacpacity. The Dyn_prev_buffer model considers Buffer_UP because it is slightly better than Buffer_DOWN. The choice to benchmark against static PM and fixed-buffer policies, rather than against CBM or RL approaches, is deliberate for three reasons. First, static policies remain the industrial baseline in many manufacturing settings due to their simplicity and low sensor requirements—

establishing improvement over this baseline is a necessary first step. Second, CBM requires continuous degradation monitoring (e.g., vibration, temperature sensors) that our proposed policy explicitly avoids, making direct comparison methodologically inconsistent. Third, RL approaches (e.g., [26]) optimize policies for specific system configurations through extensive training, whereas our objective is a heuristic that generalizes without retraining. A direct numerical comparison with these approaches would require identical assumptions about sensor availability and training budgets, which is outside this paper's scope. However, Section 6 acknowledges the integration of predictive methods as a future direction.

The simulation results across the 16 cases reveal meaningful patterns in the influence of system parameters on unit cost. The average base cost across all scenarios was 55.44, with a standard deviation of 12.67, indicating considerable sensitivity to parameter changes.

The effect of buffer size is clearly reflected in the Buffer_UP and Buffer_DOWN cost columns. Increasing buffer capacity (from 5 to 10 units) consistently led to cost reductions, with an average decrease of approximately 29% across all scenarios. For instance, Case 1 saw a 29.82% reduction when the buffer was increased. Conversely, reducing the buffer caused increases in unit cost, though less pronounced, typically in the range of 8–21% (e.g., +20.95% in Case 8). These results highlight that larger buffers improve system stability and reduce costs, especially under high load or limited preventive scheduling flexibility.

The influence of preventive maintenance frequency (interval prev) is visible in the Dyn_prev column. Longer intervals between preventive actions generally led to lower unit costs, with percentage reductions exceeding 50% in some cases (e.g., Case 1: −52.48%). However, this comes with potential risks related to reliability that may not be captured in cost metrics alone. This suggests that aggressive preventive scheduling may not always be economically optimal, particularly when other buffers or redundancies are present.

The Dyn_prev_buffer column, which reflects the combined effect of all changes, shows that the most cost-effective configurations (e.g., Case 1 and Case 9) achieved over 50% reduction in unit cost relative to the baseline. Interestingly, Cases such as 12 and 16 reported positive percentage variations (+9.66% and +13.57%, respectively), indicating that in some configurations, parameter changes may lead to worse performance. This underlines the importance of multi-factor interactions and context-specific tuning.

The findings confirm that buffer capacity and maintenance interval are the most influential levers for cost optimization. While increased buffer sizes consistently reduce cost, the effect of preventive maintenance scheduling is more nuanced and dependent on system dynamics and interaction with other parameters. These insights can guide the design of robust operational policies in similar production systems.

The Dyn_prev_buffer configuration represents the combined effect of all considered parameters—buffer capacity, preventive maintenance interval and duration, and bottleneck location. Among all performance indicators analyzed, this model consistently achieved the lowest unit costs, with an average of 39.17 and a standard deviation of 7.80, which is notably lower than in the base scenario (mean = 55.44; std = 12.67).

In particular, Case 1 and Case 9, both based on the "full variation" setting, reached the most significant cost reductions of −52.81% and −53.58%, respectively. These configurations combined longer preventive intervals, lower maintenance time, larger buffer capacity, and a bottleneck placed at 50% or 25%, highlighting how synergistic effects among parameters can lead to substantial efficiency gains.

However, it is also worth noting that cases 12 and 16 under the full variation model showed positive cost variations (+9.66% and +13.57%), despite using the same parameter variation structure. This suggests that the interaction with bottleneck location or specific parameter alignment may influence whether the full variation model achieves optimal

results. Therefore, parameter tuning must be context-specific, and a universal "best setting" may not exist across all operational conditions.

Overall, the Dyn_prev_buffer model emerges as the most promising configuration in terms of cost efficiency, but its effectiveness depends on careful coordination of all factors, especially buffer management and preventive scheduling. Future work could extend this analysis to explore interaction effects in more detail and test the robustness of this model under different demand or failure scenarios.

Then, the ANOVA was studied to evaluate the impact of the factors studied. Table 6 reports the ANOVA and regression for the case with dynamic preventive maintenance and fixed buffer capacity. All ANOVA tests were conducted at a significance level of $\alpha = 0.05$, and factors p-values below this threshold were considered statistically significant. The fixed time between preventive maintenance actions is denoted by "actions_prev" for the case of every 10 h or 30 h.

Table 6. ANOVA dynamic preventive policy with fixed buffer capacity.

Source of Variance	Sum of Squares	Degrees of Freedom	Mean of Squares	F-Value	p-Value
bottleneck	178.42	1	178.42	4.26	0.044
actions prev	1809.12	1	1809.12	43.19	0.000
Tp	65.87	1	65.87	1.57	0.215
buffer cap	2313.09	1	2313.09	55.22	0.000
threshold buffer	6.14	1	6.14	0.15	0.703
threshold time	108.29	1	108.29	2.59	0.113
Error	2387.76	57	41.89		
Total	6868.69	63			
Regression model					
costs prev = 42.951 + 1.670 * bottleneck_25.00% − 1.670 * bottleneck_50.00% − 5.317 * actions prev_10 + 5.317 * actions prev_30 + 1.015 * Tp_3 − 1.015; Tp_5 − 6.012 * buffer cap_5 + 6.012 * buffer cap_10 − 0.310 * threshold buffer_2 + 0.310 * threshold buffer_3 − 1.301 * threshold time_50.00% + 1.301 * threshold time_75.00%					

The ANOVA results indicate that actions prev ($p = 0.000$) and buffer cap ($p = 0.000$) have a highly significant effect on the dependent variable (costs prev), with F-values of 43.19 and 55.22, respectively. These factors are the most influential parameters in the model. The factor bottleneck also shows a statistically significant effect ($p = 0.044$), though with a lower F-value (4.26), suggesting a more moderate impact. Conversely, Tp, threshold buffer, and threshold time have p-values > 0.05, indicating no statistically significant effect at the 95% confidence level.

The regression equation provides insight into how each factor level influences the predicted costs. Notably:

- Increasing *actions prev* from level 10 to 30 results in a cost increase of approximately 10.63 units (−5.317 to +5.317), indicating a strong nonlinear relationship.
- Increasing *buffer cap* from level 5 to 10 leads to a significant increase in costs (+12.024), aligning with its strong significance in the ANOVA.
- The *bottleneck* variable contributes symmetrically: increasing from 25% to 50% reduces the cost effect from +1.670 to −1.670, suggesting a symmetric contrast coding structure.
- Other variables, such as *threshold buffer* and *threshold time*, have much smaller coefficients, consistent with their lack of statistical significance.

From a practical standpoint, the analysis suggests that modifying the number of preventive actions and adjusting buffer capacities are the most effective levers for reducing or controlling costs. The limited influence of threshold-related variables suggests that these may be deprioritized or simplified in decision-making models.

Table 7 reports the ANOVA and regression for the case with dynamic preventive maintenance and dynamic buffer capacity.

Table 7. ANOVA dynamic preventive policy and dynamic buffer capacity.

Source of Variance	Sum of Square	Degree of Freedom	Mean of Square	F-Value	*p*-Value
bottleneck	1164.82	1	1164.82	92.10	0.000
Tp	371.35	1	371.35	29.36	0.000
preventive time	69.43	1	69.43	5.49	0.027
threshold buffer	7.43	1	7.43	0.59	0.450
threshold time	11.75	1	11.75	0.93	0.344
error	328.81	26	12.65		
total	1953.59	31			
Regression model					
costs = 40.0299 + 6.033 * bottleneck_25.00% − 6.033 * bottleneck_50.00% − 3.407 * Tp_10 + 3.407 * Tp_30 + 1.473 * preventive time_3 − 1.473 * preventive time_5 − 0.482 * threshold buffer_2 + 0.482 * threshold buffer_3 − 0.606 * threshold time_50.00% + 0.606 * threshold time_75.00%					

The ANOVA table shows that three factors—bottleneck, Tp, and preventive time—have a statistically significant impact on the response variable (costs), with *p*-values well below the 0.05 threshold. Bottleneck stands out as the most influential factor, with an F-value of 92.10 and a very strong statistical significance ($p = 0.000$), explaining a large portion of the variance (Adj SS = 1164.82). Tp also contributes significantly ($F = 29.36$, $p = 0.000$), followed by preventive time ($F = 5.49$, $p = 0.027$), though the latter shows a more moderate effect.

On the other hand, threshold buffer and threshold time show no significant effect (*p*-values > 0.3), suggesting that these variables do not meaningfully influence the outcome in the current experimental setup.

The regression equation further clarifies the direction and magnitude of each factor's influence:

The bottleneck variable shows a symmetric contrast between levels (±6.033), reinforcing its critical role in cost variations.

For interval prev, moving from 10 to 30 increases costs by 6.814 units (−3.407 to +3.407), again highlighting its strong impact.

Tp also affects costs, though more modestly (±1.473), consistently with the ANOVA findings.

The coefficients for threshold buffer and threshold time are small and directionally symmetric (±0.482 and ±0.606, respectively), aligning with their lack of significance in the ANOVA.

These findings suggest that system bottlenecks and preventive maintenance intervals are the primary levers for cost management in the studied context. In contrast, threshold-related parameters may be less critical and could potentially be excluded or deprioritized in operational decision-making. The results also imply that minor adjustments in preventive timing (Tp) have measurable, albeit limited, impact.

Figure 3 and Table 8 report the working state of the flow line for the five models over the 16 cases.

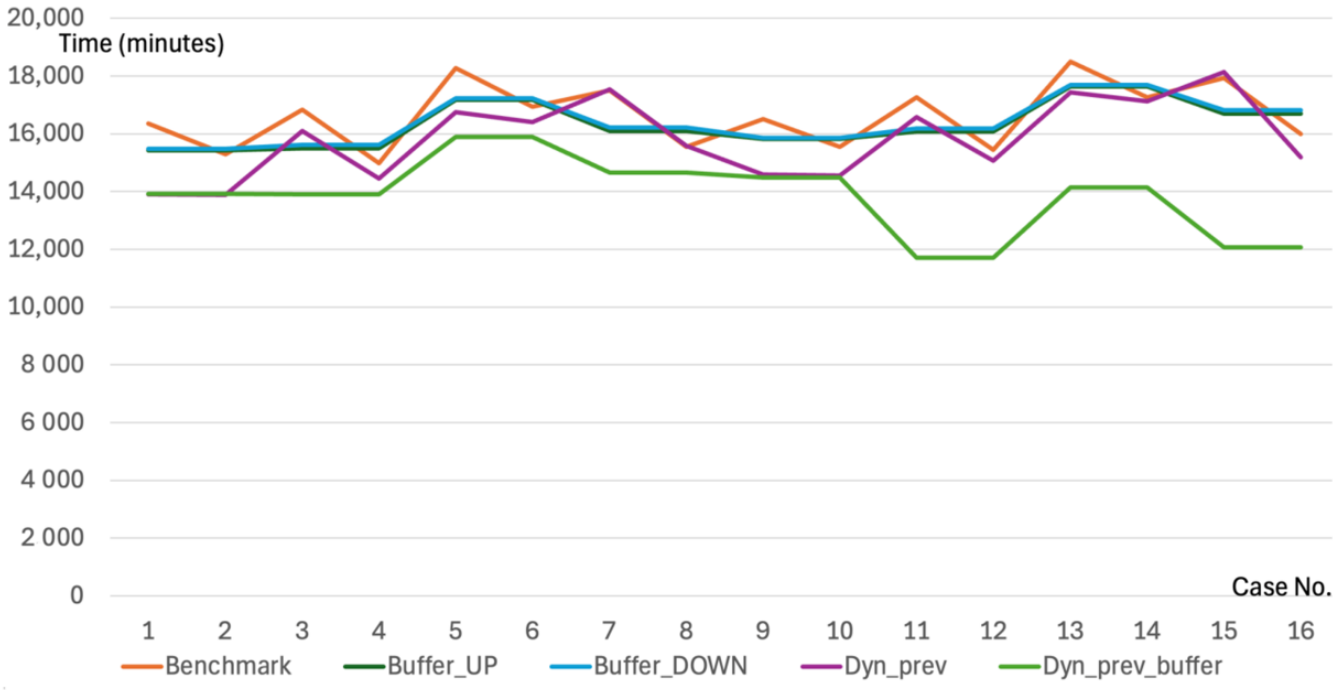


Figure 3. Working time.

Table 8. Working state of the machines.

Case No.	Buffer_UP	Buffer_DOWN	Dyn_Prev	Dyn_Prev_Buffer
1	−5.62%	−5.33%	−14.96%	−14.81%
2	0.89%	1.20%	−9.19%	−8.93%
3	−7.92%	−7.24%	−4.38%	−17.43%
4	3.43%	4.19%	−3.57%	−7.25%
5	−6.03%	−5.74%	−8.37%	−13.06%
6	1.40%	1.71%	−3.11%	−6.19%
7	−8.06%	−7.34%	0.22%	−16.30%
8	3.49%	4.30%	0.24%	−5.79%
9	−4.17%	−3.99%	−11.63%	−12.33%
10	1.74%	1.93%	−6.43%	−6.92%
11	−6.82%	−6.27%	−3.91%	−32.19%
12	4.08%	4.70%	−2.50%	−24.25%
13	−4.62%	−4.33%	−5.72%	−23.58%
14	2.24%	2.55%	−0.72%	−18.08%
15	−6.81%	−6.21%	1.16%	−32.70%
16	4.48%	5.15%	−4.98%	−24.55%
Average	−1.77%	−1.29%	−4.87%	−16.52%

The analysis of machine processing time variation across the different simulation models reveals important trends regarding the influence of system configuration changes on operational performance.

The Dyn_prev_buffer model consistently delivers the largest reductions in machine processing time, with an average decrease of −16.52%, significantly outperforming the other configurations. This suggests that joint optimization of buffer size, preventive maintenance scheduling, and bottleneck location can have a strong synergistic impact on process efficiency, leading to more fluid production flows and fewer delays due to machine downtime or accumulation.

Notably, cases 11, 12, 13, 15, and 16 all show processing time reductions exceeding −23%, with case 15 achieving the largest reduction at −32.70%. These results imply that aligning a longer preventive maintenance interval, reduced maintenance duration, strategic buffer placement, and a shifted bottleneck can significantly lower machine workload and improve time-based performance.

In contrast, the Buffer_UP and Buffer_DOWN configurations show more modest improvements, with mean processing time variations of only −1.77% and −1.29%, respectively. Interestingly, these models sometimes even result in increased processing times (e.g., Case 4 and Case 8), suggesting that isolated buffer adjustments may lead to imbalanced flow or ineffective utilization without concurrent changes to maintenance or layout strategies.

The Dyn_prev model, which adjusts only the preventive maintenance policy, shows an intermediate average improvement (−4.87%). While beneficial, this effect is considerably smaller than in the full variation scenario, indicating that maintenance policy alone is insufficient to drive significant processing efficiency gains.

In summary, these findings confirm that integrated system design is crucial for optimizing production time. The Dyn_prev_buffer model stands out not only in terms of unit cost reduction but also in minimizing machine processing time, thereby contributing to higher throughput and better resource utilization. This highlights the importance of multi-parameter tuning in production system optimization, especially when facing complex trade-offs between cost, time, and reliability.

Figure 4 reports the waiting and blocked state of the flow line for the five models over the 16 cases. The chart clearly shows that the Dyn_prev_buffer model consistently achieves the highest performance, maintaining lower peaks in waiting/blocked states for most cases while adapting effectively to fluctuating system conditions. In contrast, models such as Buffer_UP exhibit higher and more frequent peaks, indicating prolonged blocked states, which translate into reduced throughput efficiency.

The cyclical pattern observed suggests alternating phases of congestion and recovery, particularly in benchmark and Buffer_DOWN, pointing to the limitations of static control strategies. The responsiveness of dyn_prev_buffer highlights the potential of adaptive policies to anticipate congestion and balance flow dynamically.

These results suggest that a start-and-stop policy could be further integrated: by monitoring real-time buffer and blockage levels, the system could proactively pause certain stages to prevent upstream congestion and then resume operations when downstream capacity is restored. This approach would combine adaptive maintenance and buffer allocation with targeted production halts, further reducing inefficiencies and costs.

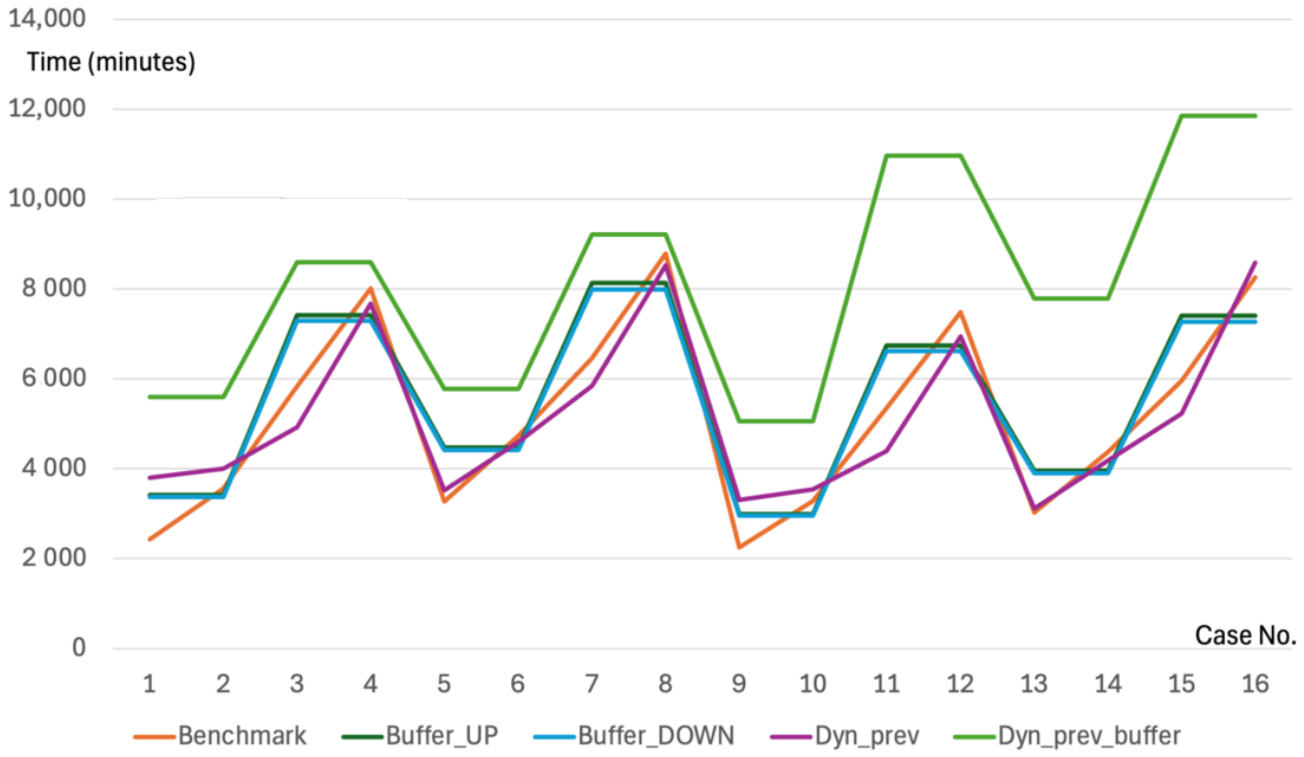


Figure 4. Waiting and blocked.

6. Conclusions and Future Development Paths

This study proposed and comprehensively evaluated an integrated, adaptive control strategy designed to address the challenges of variability and efficiency in modern production systems. The research provided novel contributions in three key areas: the design of an adaptive policy, its evaluation in a multi-product context, and the inclusion of an energy consumption proxy.

The results conclusively demonstrate the superiority of the fully integrated policy (Dyn_prev_buffer). This adaptive approach achieved remarkable performance gains, including unit cost reductions exceeding 50% in optimal scenarios and an average reduction in machine working time of 16.5% compared to the benchmark. Crucially, these results were validated across a multi-product environment simulated through shifting bottleneck locations, confirming the policy's robustness and flexibility in managing varied operational dynamics. Furthermore, the significant reduction in machine working time provides a strong proxy for enhanced energy efficiency, demonstrating that the policy not only cuts operational costs but also contributes to more sustainable manufacturing practices.

Statistical analysis confirmed that these improvements stem from the synergistic effect of dynamically coordinating maintenance and buffer levers. ANOVA identified preventive maintenance intervals and buffer capacity as the most influential control parameters, while the specific real-time thresholds for triggering maintenance were found to be not statistically significant, suggesting an opportunity to simplify control logic.

The findings of this study offer several practical insights for managers aiming to optimize manufacturing systems:

- Adopt an Integrated Approach: The results strongly suggest moving beyond the "siloed" management of maintenance and inventory policies. Jointly designing and managing these two areas can unlock efficiency gains that are unattainable by optimizing them separately.

- Focus on Key Levers: Optimization efforts should concentrate on the parameters with the greatest impact: buffer capacity and the scheduling of preventive maintenance intervals. The lack of significance from the maintenance activation thresholds suggests that companies can implement simpler, more robust control logic without compromising performance.
- Invest in Data-Driven Flexibility: The superiority of dynamic models highlights the strategic value of investing in enabling Industry 4.0 technologies (sensors, IoT, data analytics). These tools allow for a shift from static control policies to adaptive strategies that improve not just cost, but also operational robustness and energy efficiency.

While this study provides valuable insights, it is subject to certain limitations that open several avenues for future research. First, the proposed adaptive strategy was tested on a serial flow line; future work should validate these policies on more complex manufacturing layouts, such as job shops or flexible manufacturing systems. Second, the current model assumes exponential failure distributions. Incorporating more realistic degradation models, such as the Weibull distribution, would better capture actual component wear over time. Furthermore, a significant step forward involves integrating Machine Learning (ML) for predictive maintenance. By employing Prognostics and Health Management (PHM) techniques, future models could estimate a machine's Remaining Useful Life (RUL) in real time, shifting from a flow-aware heuristic to a fully condition-based predictive control system. A limitation of this study is the absence of direct comparison with CBM or RL-based policies. While such comparison is methodologically challenging due to different information requirements (sensors for CBM; training episodes for RL), future work should establish relative performance under standardized assumptions. The use of working time as an energy proxy is a clear limitation. Actual energy consumption depends on machine-specific power curves, with idle states consuming non-zero energy and partial loads exhibiting non-linear efficiency. Future work should incorporate physical energy models or actual power measurements. Finally, applying formal multi-objective optimization would explicitly analyze the trade-offs (Pareto front) between cost, throughput, and resilience, while simultaneously testing the policy's robustness against variable demand and supply chain disruptions.

Supplementary Materials: The raw data supporting the conclusions of this article: https://www.mdpi.com/article/10.3390/app16104607/s1.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The raw data supporting the conclusions of this article is available in Supplementary File.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Kagermann, H.; Wahlster, W. Ten Years of Industrie 4.0. *Sci* **2022**, *4*, 26. [CrossRef]
2. Kuhnle, A.; Kaiser, J.-P.; Theiß, F.; Stricker, N.; Lanza, G. Designing an adaptive production control system using reinforcement learning. *J. Intell. Manuf.* **2021**, *32*, 855–876. [CrossRef]
3. Jain, A.; Bhatti, R.; Singh, H. Total productive maintenance (TPM) implementation practice: A literature review and directions. *Int. J. Lean Six Sigma* **2014**, *5*, 293–323. [CrossRef]
4. Hopp, W.J.; Spearman, M.L. *Factory Physics 2011*; Waveland Press: Long Grove, IL, USA, 2011.

5. Wang, F.; Lu, Y.; Ju, F. Condition-based real-time production control for smart manufacturing systems. In Proceedings of the 2018 IEEE 14th International Conference on Automation Science and Engineering (CASE), Munich, Germany, 20–24 August 2018; pp. 1052–1057.
6. Zhang, N.; Si, W. Deep reinforcement learning for condition-based maintenance planning of multi-component systems under dependent competing risks. *Reliab. Eng. Syst. Saf.* **2020**, *203*, 107094. [CrossRef]
7. Renna, P. Adaptive policy of buffer allocation and preventive maintenance actions in unreliable production lines. *J. Ind. Eng. Int.* **2019**, *15*, 411–421. [CrossRef]
8. Gaber, Y.H.; El-Khodary, I.A.; Abdelsalam, H.M. Joint optimization of part quality inspection planning, buffer allocation, and preventive maintenance in a serial manufacturing. *J. Adv. Manuf. Syst.* **2025**, *24*, 435–469. [CrossRef]
9. Lee, J.; Bagheri, B.; Kao, H.-A. A cyber-physical systems architecture for industry 4.0-based manufacturing systems. *Manuf. Lett.* **2015**, *3*, 18–23. [CrossRef]
10. Ye, Y.; Grossmann, I.E.; Pinto, J.M.; Ramaswamy, S. MINLP model for reliability optimization of system design and maintenance based on markov chain representation. *Comput. Aided Chem. Eng.* **2020**, *48*, 1057–1062.
11. Tsarouhas, P.H. Performance evaluation of the croissant production line with reparable machines. *J. Ind. Eng. Int.* **2015**, *11*, 101–110. [CrossRef]
12. Rao, M.S.; Naikan, V.N.A. Reliability analysis of repairable systems using system dynamics modeling and simulation. *J. Ind. Eng. Int.* **2014**, *10*, 69. [CrossRef]
13. Ferreira, C.; Gonçalves, G. Remaining Useful Life prediction and challenges: A literature review on the use of Machine Learning Methods. *J. Manuf. Syst.* **2022**, *63*, 550–562. [CrossRef]
14. Dui, H.; Yang, X.; Fang, Y. Evaluation methodology for preventive maintenance in multi-state manufacturing systems considering different costs. *Int. J. Prod. Res.* **2023**, *61*, 8309–8324. [CrossRef]
15. Guner, H.U.; Murat, A.; Chinnam, R.B. Performance analysis of a production system with two deteriorating machines and a finite buffer. *Cent. Eur. J. Oper. Res.* **2024**, *32*, 621–651. [CrossRef]
16. Magnanini, M.C.; Terkaj, W.; Tolio, T.A.M. A novel algorithm for optimal buffer allocation in automated asynchronous unreliable lines. *Int. J. Adv. Manuf. Technol.* **2022**, *122*, 4405–4419. [CrossRef]
17. Neto, W.A.F.; Cavalcante, C.A.V.; Do, P. Deep reinforcement learning for maintenance optimization of a scrap-based steel production line. *Reliab. Eng. Syst. Saf.* **2024**, *249*, 110199. [CrossRef]
18. Weiss, S.; Schwarz, J.A.; Stolletz, R. The buffer allocation problem in production lines: Formulations, solution methods, and instances. *IISE Trans.* **2019**, *51*, 456–485. [CrossRef]
19. Dolgui, A.; Eremeev, A.V.; Sigaev, V.S. Fitness landscapes of buffer allocation problem for production lines with unreliable machines. *Comput. Oper. Res.* **2024**, *172*, 106819. [CrossRef]
20. Zhao, Z.; Goodall, P.; West, A.; Colligan, A.; Friel, I.; Hickinbotham, S.; Price, M.; Jin, Y. A bio-inspired evolution-development method for modelling and optimisation of buffer allocation in unreliable serial production line. In Proceedings of the 3rd International Conference on Mechanical, Aerospace and Automotive Engineering (CMAAE 2023), Nanjing, China, 8–10 December 2023; Volume 2023, pp. 90–95.
21. Li, L.; Dong, M.; Zhou, J. A Recursive Optimization Approach for Buffer Allocation in Large Production Lines. *Appl. Sci.* **2024**, *14*, 9953. [CrossRef]
22. Li, L.; Luo, X.; Zhang, S.-G. A scale transformation algorithm for buffer allocation in production lines with unreliable machines. *Comput. Ind. Eng.* **2025**, *201*, 110918. [CrossRef]
23. Gaber, Y.H.; El-Khodary, I.A.; Abdelsalam, H.M. A model review on joint optimization of part quality inspection planning, buffer allocation, and preventive maintenance in SMMS. *J. Adv. Manuf. Syst.* **2023**, *22*, 667–691. [CrossRef]
24. Nahas, N. Buffer allocation and preventive maintenance optimization in unreliable production lines. *J. Intell. Manuf.* **2017**, *28*, 85–93. [CrossRef]
25. Motlagh, M.M.; Azimi, P.; Amiri, M.; Madraki, G. An efficient simulation optimization methodology to solve a multi-objective problem in unreliable unbalanced production lines. *Expert Syst. Appl.* **2019**, *138*, 112836. [CrossRef]
26. Hung, Y.-H.; Shen, H.-Y.; Lee, C.-Y. Deep reinforcement learning-based preventive maintenance for repairable machines with deterioration in a flow line system. *Ann. Oper. Res.* **2024**, 1–21. [CrossRef]
27. Zandieh, M.; Joreir-Ahmadi, M.N.; Fadaei-Rafsanjani, A. Buffer allocation problem and preventive maintenance planning in non-homogenous unreliable production lines. *Int. J. Adv. Manuf. Technol.* **2017**, *91*, 2581–2593. [CrossRef]
28. Bányai, Á.; Bányai, T. Real-time maintenance policy optimization in manufacturing systems: An energy efficiency and emission-based approach. *Sustainability* **2022**, *14*, 10725. [CrossRef]
29. Salmasnia, A.; Shabani, A. Opportunistic maintenance modeling for series production systems based on bottleneck by considering energy consumption and market demand. *J. Ind. Prod. Eng.* **2023**, *40*, 506–518. [CrossRef]
30. Pinciroli, L.; Baraldi, P.; Zio, E. Maintenance optimization in industry 4.0. *Reliab. Eng. Syst. Saf.* **2023**, *234*, 109204. [CrossRef]

31. Elsayed, E.A. Reliability, Maintainability, Safety, and Sustainability. In *Springer Handbook of Automation*; Springer International Publishing: Cham, Switzerland, 2023; pp. 699–715.
32. Meller, R.D.; Kim, D.S. The impact of preventive maintenance on system cost and buffer size. *Eur. J. Oper. Res.* **1996**, *95*, 577–591. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.