

Article

Towards Intelligent Fused Filament Fabrication: Computational Verification of a Monitoring and Early-Warning Framework for Instability Mitigation

Massimo Pacella ^{1,*} , Antonio Papa ²  and Gabriele Papadia ¹ 

¹ Department of Engineering for Innovation, University of Salento, 73100 Lecce, Italy; gabriele.papadia@unisalento.it

² Department of Science and Information Technology, Pegaso University, 80121 Napoli, Italy; antonio.papa@unipegaso.it

* Correspondence: massimo.pacella@unisalento.it

Abstract

Fused Filament Fabrication (FFF) plays a critical role in several application fields due to its affordability and manufacturing versatility. However, FFF reliability remains vulnerable to rapid environmental and operational variations, which directly influence the dimensional precision and mechanical properties of printed parts. To address these challenges, this study presents a simulation-based computational framework for the real-time early-warning supervision of FFF systems. The proposed multilayer architecture integrates high-throughput data acquisition, distributed computing, and dynamic analysis to proactively detect deviations from optimal conditions. Architectural verification follows a simulation-first methodology designed to replicate the operational dynamics of standard FFF hardware. By employing telemetry streams to test the decision-making pipeline, the study isolates computational performance, such as throughput and latency, from the confounding variables of physical hardware. This approach enables a precise, deterministic assessment of the system's responsiveness, serving as a foundational de-risking step prior to empirical implementation. Numerical results of this study show that the integrated distributed computing model successfully manages high-frequency telemetry with a response time within the operational safety margins, confirming the architectural viability of the proposed solution. By providing insights into system behavior prior to physical deployment, this simulation-first strategy mitigates implementation risks and offers practical guidance for developing autonomous additive manufacturing workflows, advancing the transition toward intelligent industrial FFF.

Keywords: fused filament fabrication; distributed computing; simulation framework; virtual verification; control architecture



Academic Editor: Abílio M. P. De Jesus

Received: 28 November 2025

Revised: 27 December 2025

Accepted: 29 December 2025

Published: 1 January 2026

Copyright: © 2026 by the authors.

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](https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The integration of the Internet of Things (IoT) and Cloud Computing has accelerated the evolution of Cloud Manufacturing (CM), redefining the connection, monitoring, and optimization of distributed production resources. Currently, CM represents the central pillar of the Industry 4.0 paradigm, enabling production environments driven by real-time decision-making and continuous feedback. In CM, the use of simulation, automation, and data analytics is evolving production processes into interconnected and self-adaptive networks capable of proactive supervision and autonomous optimization.

Additive Manufacturing (AM), fundamentally defined by the layer-by-layer material consolidation principle pioneered by Hull [1], has evolved from a rapid prototyping tool to a strategic production technology. Characterized by layer-by-layer fabrication of complex geometries, AM ensures material efficiency, design flexibility, and reduced production lead times [2]. According to ASTM standards [3,4], material extrusion processes, particularly Fused Filament Fabrication (FFF), remain among the most accessible and versatile AM technologies, with applications in the aerospace, biomedical, and automotive sectors [5].

As shown in Figure 1, FFF operates by feeding a thermoplastic filament into a heated nozzle, melting it, and depositing successive layers to build the printed part. The quality of the final component depends on both machine and environmental variables, such as extrusion temperature, printing speed, and ambient humidity [6]. However, the quality of FFF parts is highly sensitive to a multitude of process variables. As extensively documented in literature, mechanical properties and dimensional accuracy are strictly dependent on parameters such as layer thickness, infill density, and raster angle [7]. While these factors induce a characteristic anisotropy in the printed components [8], the stability of the deposition process itself is the primary prerequisite for structural integrity. Variations in these parameters can lead to defects such as warping, delamination, and porosity, which compromise dimensional accuracy and mechanical performance [9,10].




Figure 1. Key components of the FFF process. A thermoplastic filament is fed into a heated nozzle via a gear system, melted, and deposited layer by layer onto a build plate to form the final part.

Polylactic Acid (PLA) is one of the most widely used thermoplastics in FFF. However, this material exhibits critical dependencies on environmental conditions owing to its hygroscopic and temperature-dependent viscoelastic behaviors [11,12]. For example, a relative humidity (RH) exceeding 50% results in accelerated moisture uptake by PLA, causing several undesired phenomena such as hydrolytic degradation, molecular weight reduction, and altered rheological properties [13,14]. These phenomena may result in poor interlayer adhesion, melt instability, and dimensional deformation [15,16]. Critically, standard open-loop FFF controllers remain blind to these fluctuations, continuing to operate under the assumption of constant material properties despite active degradation. This ultimately impacts the sustainability of the FFF process through increased waste, energy consumption, and reworking [17].

Therefore, ensuring repeatable and reliable outcomes in extrusion-based AM requires effective monitoring and control of environmental and process conditions. Establishing robust relationships between process parameters and part quality typically requires large experimental datasets, the acquisition of which can be time-consuming and expensive.

Computational modeling and simulation offer efficient and reproducible alternatives, enabling the systematic exploration of parameter interactions and verification of control logic stability under diverse operating conditions without the material constraints of physical experiments.

Despite these advancements, a critical scientific gap persists regarding the integration of material-aware constraints into real-time monitoring pipelines. Although IoT-based data logging and material degradation have been explored independently, these domains remain largely disconnected in practical architectures. Current supervisory systems often lack the computational responsiveness and integration necessary to dynamically couple fluctuating environmental conditions, such as RH, with the physical constraints of the polymer state. This fragmentation results in “blind monitoring” architectures: systems that log sensor deviations without contextualizing their impact on material integrity, thereby limiting the ability to identify process instabilities at the data-ingestion stage before they escalate into irreversible structural flaws.

In this study, we propose a simulation-based computational framework for the early-warning monitoring and supervision of FFF systems. The architecture comprises three main functional layers: (i) a high-throughput data ingestion layer for capturing environmental and machine variables, (ii) a distributed processing layer for real-time computation and analysis, and (iii) an application layer for anomaly detection and preventive warnings. Using simulated data streams, the framework allows for the controlled evaluation of system responsiveness, latency, and scalability before physical deployment.

This simulation-driven verification represents a foundational de-risking step toward full implementation. By strictly isolating architectural performance, specifically throughput and decision latency, from the confounding variables of physical hardware, this study provides a deterministic assessment of the system’s operational readiness. The architecture supports integration with real FFF printers to enable continuous monitoring of key parameters such as nozzle temperature, layer thickness, and printing speed. Our research contributes to the applied study of AM technologies by presenting a reproducible, scalable, and data-driven approach for monitoring and optimization in FFF, ultimately reducing the material and energy costs associated with trial-and-error hardware validation.

The remainder of this paper is structured as follows. Section 2 reviews current monitoring strategies and related work; Section 3 describes the main contributions of the proposed framework; Section 4 details the methodology and computational verification setup; Section 5 presents the computational performance results; and Section 6 discusses the implications and future research directions. The conclusions are drawn in the final section, and Appendices A and B provide numerical results and the implementation of the proposed algorithms.

2. Literature Review

The workflow of an FFF system typically comprises Computer-Aided Design (CAD), STL generation, slicing, and physical deposition. Despite the widespread use of FFF systems for rapid prototyping in aerospace, biomedical, and automotive fields [18], the technology still faces challenges regarding dimensional accuracy, process repeatability, and the structural integrity of printed components.

Machine and environmental parameters play decisive roles in determining process stability. Lendvai et al. [9] showed that even small variations in ambient temperature and relative humidity can significantly alter the rheological characteristics of polymer melts, affecting extrusion consistency and interlayer bonding. These fluctuations often lead to defects such as warping, delamination, and internal porosity. PLA, a standard thermoplastic, is particularly susceptible to humidity due to its hygroscopic and viscoelastic

nature. When exposed to humidity levels above 50%, PLA absorbs moisture, resulting in hydrolytic degradation. This process reduces molecular weight and decreases mechanical performance [13,14,16]. The resulting structural weaknesses, characterized by lower tensile strength and weakened interlayer adhesion [15,19], highlight environmental conditioning as a key factor for process robustness and product quality [17].

While prior studies have provided insights into environmental effects on material performance, the literature remains fragmented regarding the early warning detection of anomalies induced by fluctuating ambient conditions. Specifically, the coupling between environmental variability, filament rheology, and extrusion stability is rarely captured by existing monitoring systems. This limitation hinders the implementation of proactive supervisory architectures capable of identifying anomalies—such as unstable flow, nozzle clogging, and layer deformation—before they impact the final product. To address this gap, the data-rich nature of AM processes supports the integration of real-time sensing and advanced analytics, facilitating anticipatory anomaly detection and adaptive control.

2.1. Literature Search Strategy

To systematically assess existing contributions to environmental monitoring, process optimization, and computational analytics in extrusion-based AM, we conducted a structured literature review. The selection process was guided by the principles of the PRISMA statement [20], ensuring transparency and reproducibility without the extensive reporting requirements of a standalone systematic review.

The methodology followed a sequential filtering logic comprising identification, screening, and eligibility assessment. Relevant literature was retrieved from major engineering databases, specifically Scopus and ScienceDirect, using predefined search strings and Boolean operators. To ensure relevance of the reviewed studies, strict eligibility criteria were applied. Table 1 summarizes the inclusion and exclusion parameters used to filter the initial dataset, targeting peer-reviewed developments from 2018 to 2025 that address monitoring, IoT architectures, or environmental effects in polymer AM. Consequently, Metal AM was excluded, as it entails distinct physical dynamics unrelated to polymer extrusion.

Table 1. Inclusion and exclusion criteria for the literature selection process.

Criterion	Inclusion Criteria	Exclusion Criteria
Time Period	2018–2025	Pre-2018 studies (unless foundational)
Language	English	Non-English publications
Document Type	Peer-reviewed journal articles, book chapters	Conference abstracts, editorials, patents
Topic Relevance	FFF/FDM monitoring, IoT/cloud architectures in AM, environmental effects on polymers	Metal-based AM, purely structural analysis without monitoring focus

As illustrated in the PRISMA flow diagram (Figure 2), the initial search yielded 95 records. After removing duplicates, the articles underwent a multi-stage screening based on title, abstract, and full-text relevance. Ultimately, 25 articles were selected for the final qualitative synthesis, confirming their alignment with the research objectives regarding sensing integration and data-driven supervision and distinguishing between physical experimentation and computational verification approaches.

PRISMA Flow diagram for identifying research related to the topic of environmental monitoring, process optimization and computational analytics in extrusion-based AM

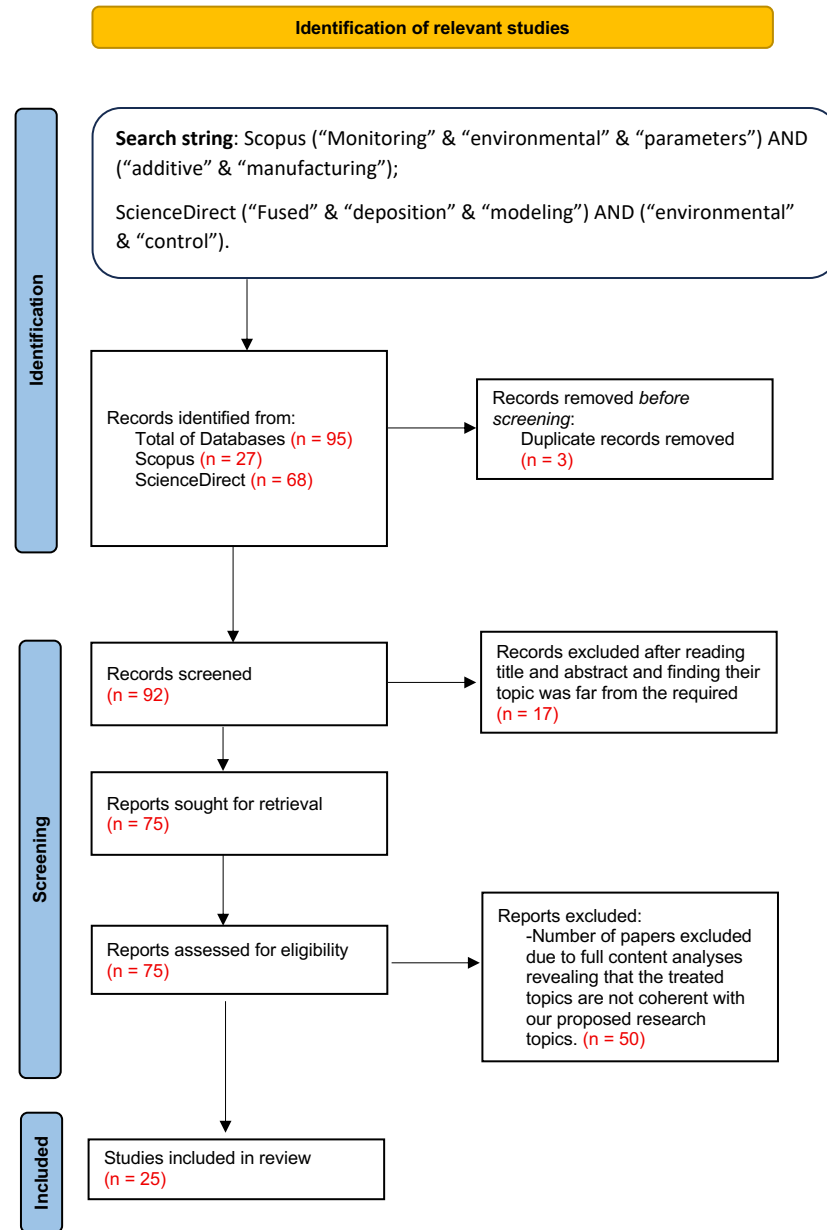


Figure 2. PRISMA flow diagram outlining the systematic literature search and selection process.

2.2. Sensing Architectures and IoT Integration in AM

The literature review identified a distinct group of studies centered on the convergence of sensing technologies, communication protocols, and computational infrastructure for process supervision. These studies emphasize the shift from isolated measurements to integrated IoT architectures.

Mazzarisi et al. [21] implemented infrared thermography for real-time melt pool observation in directed energy deposition, whereas Borghetti et al. [22] proposed embedded sensor architectures within polymer components for in situ temperature and strain tracking. Fedullo et al. [23] examined IoT communication frameworks, such as LoRaWAN, for scalable data acquisition, and Horr et al. [24] introduced digital twin models bridging the virtual and physical domains to enable dynamic reasoning and control.

Recent advances in artificial intelligence have further enhanced defect recognition capabilities by employing convolutional neural networks and gradient boosting for surface defect and porosity analyses [25–27]. However, the integration of environmental monitoring and risk-oriented analytics remains limited, revealing the need for holistic simulation-supported architectures, such as the one proposed in this study, which unifies sensing, data streaming, and anomaly detection in a single computational environment.

2.3. Data-Driven Modeling and Material Degradation

Multiple studies indicate a significant correlation between polymer chemistry, ambient factors, and final part performance [9,10,15,19]. For PLA, humidity levels above 50% typically induce hydrolytic degradation, leading to reductions of up to 30–70% in melt viscosity and mechanical strength [14,16]. These findings confirm the necessity of continuous environmental monitoring, particularly for hygroscopic polymers, to maintain print quality and energy efficiency. Although low-cost sensor networks have been proposed [28], their effective integration with robust computational frameworks capable of low-latency decision making remains to be achieved.

Parallel developments in computational modeling and hybrid digital frameworks further reinforce the role of data-driven process optimization. Hybrid physics–data models and causal inference approaches have been employed for sustainability analyses and anomaly detection [29–31]. Machine learning techniques, such as convolutional and multi-modal designs, have shown promising capabilities in defect classification and automated diagnostics [32–36]. Despite these advancements, the absence of systems that integrate environmental sensing, virtual verification, and early-warning analytics highlights an ongoing research gap.

Building upon our prior work [37], this study addresses this gap by proposing a multilayered computational framework for FFF process monitoring. While the verification strategy employs a simulation environment to ensure algorithmic stability, the architecture is designed to align with the applied principles of Industry 4.0, focusing on digital integration, sustainability, and operational readiness validation [38].

3. Research Innovations and Contributions

Building upon the conceptual basis of the previously established CM framework [37], this study refines the technical scope to a targeted contribution: the development and computational verification of a low-latency, simulation-driven early-warning framework for FFF. The revised system is engineered to support the monitoring of pre-print and in-process environmental conditions, delivering actionable alerts whenever stability thresholds are exceeded.

The primary advancement over existing approaches lies in the tight coupling of high-frequency environmental sensing, process-specific analytics, and decision logic within a unified architecture. Unlike conventional FFF monitoring systems, which typically function as passive data loggers or post-process diagnostic tools [39,40], the proposed framework operates as an active supervisory layer. It continuously evaluates process readiness before and during printing, thereby minimizing decision latency and preventing the propagation of failure modes. By shifting the paradigm from post-process inspection to in-process risk anticipation, this approach aligns with sustainability-oriented AM practices, actively mitigating the material waste and energy expenditure associated with irreversible print failures.

At the architectural level, the general-purpose CM framework introduced in [37] was reconfigured into an FFF-specific early warning platform. Specifically, the data ingestion schemas and stream processing logic were optimized to align with the characteristic thermal time constants of the extrusion process. This reconfiguration emphasizes parameters known to critically affect extrusion stability and part quality, such as nozzle temperature, printing

speed, and ambient humidity. By embedding process-specific physical constraints directly into the computational workflow, the system assesses incoming data against dynamically defined operating windows, allowing for the prompt detection of abnormal trends rather than relying solely on static limit checks.

A distinguishing feature of this work is the adoption of a representative simulation environment for verification. Consistent with Digital Twin methodologies in AM [41], this study prioritizes architectural validation as a necessary precursor to physical integration. Rather than immediately targeting physical deployment, the use of simulation allows for the rigorous stress-testing of responsiveness and detection performance under controlled disturbances. This approach isolates computational metrics—such as latency and throughput—from the confounding factors of hardware variability, serving as a robust de-risking step prior to experimental implementation.

Material behavior is integrated at a functional level to enhance decision-making. Specifically, humidity-dependent degradation effects in PLA are abstracted into process-relevant stability indicators. While distinct from full multiphysics modeling, this logic ensures that environmental measurements are interpreted based on their practical impact on flow consistency and interlayer bonding. Consequently, alerts are triggered by anticipated process instabilities rather than raw sensor deviations alone.

Finally, the system's decision logic is explicitly designed to support graded responses. Depending on the severity and persistence of the detected anomalies, the framework can issue warnings to operators, suggest corrective actions, or signal the need for print interruption. By hierarchically categorizing process risks, this approach avoids unnecessary downtime for transient fluctuations while ensuring immediate intervention for critical failures. This dynamic adaptability distinguishes the proposed method from existing monitoring solutions that lack a formalized pathway from sensing to actionable control decisions.

4. Materials and Methods

This section outlines the methodological framework adopted for the design, verification, and performance benchmarking of the proposed architecture for early-warning monitoring. This approach was implemented in a reproducible computational environment using Python-based simulation and analytics modules.

A deterministic seed-controlled data generation routine ensured statistical traceability and formed a reliable basis for algorithmic benchmarking. The system is conceived as a proactive supervision platform; therefore, this study prioritizes the verification of the computational throughput, latency, and detection logic prior to physical deployment.

4.1. Scope and Heuristic Process–Material Integration

The framework is intentionally framed at the computational and algorithmic level, focusing on data throughput, latency, and responsiveness, rather than on the development of explicit physical–mechanical constitutive models.

To support this integration without exceeding the scope of a computational study, we implemented a rule-based logic derived from material science literature. Instead of simulating the chemical kinetics of hydrolysis or solving coupled differential equations—which would require prohibitive computational resources for a real-time supervision test—we adopted a phenomenological approach known as ‘heuristic mapping.’ This method defines specific ‘risk zones’ for PLA processing, allowing the data pipeline to process signals that reflect the logical constraints of the material:

- **Moisture Sensitivity Logic:** Literature indicates that PLA undergoes significant degradation when exposed to humidity. Based on established thresholds, RH values exceeding 50% are tagged as ‘Critical.’ In the simulation, prolonged persistence in this zone

triggers a synthetic ‘degradation drift’ in the data stream, mimicking the statistical signature of material quality loss (e.g., bubbling or inconsistent extrusion).

- **Thermal Window Logic:** The system enforces a valid viscoelastic processing window (e.g., 190–210 °C). The data points generated outside this range were correlated with simulated flow instabilities. This ensures that the monitoring logic is validated against data patterns that physically represent the onset of warping or clogging, rather than simple random noise.

4.1.1. Multivariate Conditional Logic

To address the requirement for variable correlation, without the computational overhead of full rheological modeling, we implemented a ‘logic-based coupling’ mechanism. The simulation avoids generating Temperature and Flow Rate as independent stochastic variables. Instead, a heuristic dependency rule is applied:

If the generated Nozzle Temperature falls below a critical flow threshold (e.g., 200 °C), the Printing Speed and resulting extrusion consistency are automatically attenuated via a proportional penalty factor. This simple but effective correlation mechanism allows us to verify that the monitoring system can correctly identify multivariate anomalies where a thermal failure (cause) induces a flow deviation (effect), successfully mimicking the signal signature of a ‘cold extrusion’ or nozzle clogging event.

4.1.2. Synthetic Injection of Anomaly Patterns

To ensure the monitoring algorithms are tested against realistic data behaviors rather than simple random noise, the simulation engine injects structured anomaly patterns derived from the heuristic logic defined in Section 4.1:

- **Moisture-Induced Drift Pattern:** To represent the risk of hydrolysis/delamination, the generator injects a continuous downward trend in the extrusion stability metric whenever the simulated RH exceeds the 50% threshold. This tests the system’s ability to detect long-term signal drifts.
- **Thermal Instability Pattern:** To emulate the conditions leading to warping, rapid stochastic fluctuations are superimposed on the temperature data stream. These fluctuations are programmed to violate the defined stability window, testing the latency of the ‘warning’ flags in the analytics layer.

4.1.3. Parameter Selection and Classification

In this study, the monitored variables were intentionally categorized into Ambient (environmental) and Process (machine-specific) parameters to reflect realistic FFF operating conditions. Ambient parameters include ambient temperature and relative humidity, which characterize the external exposure of the system and are known to fluctuate in non-climate-controlled environments. The process parameters comprise the nozzle temperature, printing speed, layer thickness, and infill density, which directly govern the material state, extrusion stability, and geometric fidelity of the printed part. These parameters were selected to emulate a production scenario in which environmental disturbances interact with the machine dynamics, requiring distinct supervision logic. Table 2 details this classification and the specific rationale for monitoring each variable based on its impact on part quality.

Table 2. Classification of monitored parameters and rationale for their selection.

Category	Parameter	Symbol	Rationale for Selection
Ambient	Ambient Temperature	T_{amb}	Rapid cooling drafts cause thermal contraction and corner lifting.
	Relative Humidity	RH	High moisture leads to PLA hydrolysis, causing delamination and bubbles.
Process	Nozzle Temperature	T_{nozzle}	Deviations affect viscosity. Low T causes under-extrusion; high T causes stringing.
	Printing Speed	V_{print}	High speed mismatch with flow causes under/over-extrusion.
	Layer Thickness	h_{layer}	Affects heat dissipation and surface quality.
	Infill Density	ρ_{infill}	Determines the required volumetric flow rate; deviations imply structural weakness.

4.1.4. Filament Preparation and Pre-Processing Scope

The filament preparation stage, including active drying through controlled heating cycles prior to printing, was intentionally not modeled in the present study. This decision reflects the simulation-based scope of our work, which focuses on assessing the computational feasibility, responsiveness, and architectural performance of the proposed layered monitoring framework rather than capturing the full sequence of material preprocessing steps. While filament drying is a critical and specialized operation for hygroscopic materials such as PLA, explicitly modeling drying profiles, residence times, and moisture diffusion kinetics requires detailed experimental characterization and high-fidelity physical models that are beyond the scope of this initial investigation. These aspects are currently being addressed in our ongoing experimental campaigns, where different drying cycles, temperature profiles, and ambient RH conditions are systematically evaluated to quantify their influence on the filament condition and printed part quality. The outcomes of these experiments will enable future extensions of the framework to incorporate filament preparation as an explicit upstream process variable within a fully integrated cyber–physical monitoring and control loop.

4.1.5. Parameter Selection and Model Parsimony

The monitoring platform’s scope was deliberately limited to a representative set of process and environmental variables (Nozzle Temperature, Printing Speed, Ambient Temperature, RH) to focus on validating architectural metrics like latency and throughput. This strategy prevents over-parameterization and ensures that the computational backbone is rigorously tested under controllable and reproducible disturbance conditions. Specific variables were excluded based on their signal dynamics in relation to the control horizon. For instance, the build plate temperature was not monitored dynamically, as PLA exhibits low thermal shrinkage, and bed heating typically operates as a static, closed-loop state rather than a high-frequency disturbance source [42]. Similarly, ambient illumination was omitted because photo-oxidative degradation operates on timescales exceeding the active print duration, whereas its immediate thermal effects are captured by the ambient temperature sensors. This parsimony allows for a focused evaluation of the system’s responsiveness to immediate rheological instabilities, while the modular architecture remains natively extensible to include additional sensor modalities in future cyber–physical implementations.

4.2. Architecture of the Computational Framework

The proposed architecture was designed according to a three-layer logic that emulates real-world manufacturing data flows. Figure 3 illustrates this hierarchical organization and the specific technologies selected to benchmark each layer. To ensure simulation fidelity relevant to supervisory control, the framework operates within a virtualized environment constrained by the physical limitations and data rates of commercial desktop FFF printers.

- **Data Ingestion Layer:** responsible for acquiring synthetic process and environmental variables that reproduce IoT-like sensor streams. The simulated features included ambient temperature, RH, nozzle temperature, infill density, layer thickness, and print speed.
- **Processing and Analytics Layer:** dedicated to data stream analysis, statistical aggregation, and anomaly detection. This layer identifies deviations from the nominal conditions, particularly RH and temperature fluctuations, which affect thermoplastic rheology and print quality.
- **Application Layer:** acts as the user-facing supervisory interface, visualizing monitored signals and issuing early-warning alerts whenever operational limits are exceeded, enabling timely operator response.

Humidity dynamics were prioritized because of their strong influence on the process stability during thermoplastic extrusion. Previous studies have established that PLA degradation accelerates above approximately 50% RH [9,14]. This threshold was adopted as the reference point for sensitivity assessment in the proposed framework.

PLA is known to undergo time-dependent changes in its physical and mechanical properties after extrusion, including those associated with thermal and hydrothermal aging [43]. Empirical studies [44] have shown that prolonged exposure to ambient conditions can alter PLA’s tensile strength, elastic modulus, and ductility, with statistically significant reductions in the ultimate tensile stress and elongation after accelerated aging periods of weeks to months in controlled environments. Aging processes, including chain rearrangements and hydrolytic degradation, can increase material brittleness and elevate the risk of delamination, particularly when interlayer adhesion is compromised during the printing process. Although these effects are well documented for 3D-printed PLA, the current study deliberately abstracts the material behavior to a functional level, focusing on demonstrating the feasibility and performance of the proposed layered monitoring framework under realistic environmental variability. Explicit modeling of filament aging, interlayer adhesion degradation, and other coupled physical phenomena lies beyond the present simulation scope and will be addressed in future work by integrating experimentally obtained material response data and high-fidelity physical models that capture the influence of aging on part quality.

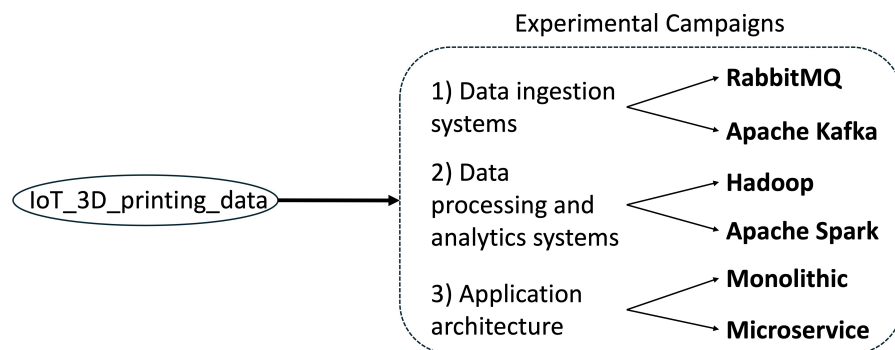


Figure 3. Overview of the simulated experimental campaign. Each layer of the proposed framework was tested independently to assess scalability, throughput, and anomaly detection responsiveness (our elaboration).

Designed as a robust digital twin precursor, the architecture currently operates on high-fidelity simulated streams to verify the logic stability. Its modular interfaces are natively ready for future integration with physical sensors (e.g., BME280 modules), cloud dashboards, and adaptive controllers for intelligent and sustainable manufacturing processes.

4.3. Simulation Dataset Generation and Reproducibility

A synthetic dataset was generated to emulate the temporal evolution of key machine and environmental parameters during FFF operation, with the primary objective of evaluating monitoring latency, data throughput, and anomaly detection performance. The simulation environment was not intended to reproduce the full thermo-mechanical physics of extrusion. Instead, it was designed to reflect realistic operational envelopes, data rates, and disturbance patterns relevant to FFF monitoring.

All simulation experiments were conducted on a workstation featuring an Intel Core™ (Tokyo, Japan) i7-11800H CPU @ 2.30 GHz and 16 GB RAM. Data generation, messaging, and analytics pipelines were implemented using Python 3.10 and the Flask framework. Locally, the distributed components were deployed using Apache Kafka v3.1.0 for message streaming and Apache Spark v3.2.1 for real-time analytics. Docker containers (v20.10) encapsulated all the services and dependencies, ensuring environmental consistency and facilitating reproducibility across various computing platforms.

Deterministic reproducibility was ensured by applying a fixed random seed across Python's random and NumPy libraries. When a previously generated dataset was detected, it was reloaded to preserve traceability; otherwise, a new dataset was generated using identical probabilistic distributions. This seed-controlled strategy ensures both reproducibility (bitwise-identical results) and repeatability, essential for robust benchmarking.

Simulation fidelity was addressed by constraining variables to operating ranges derived from Craftbot Plus printer (Craftbot, Budapest, Hungary) specifications and FFF literature [9]. To benchmark the computational throughput, a representative subset of the parameters defined in Section 4.1.3 was used. The specific bounds were selected to encompass both nominal conditions and plausible extreme scenarios:

- Ambient Temperature (T_{amb}): 20–35 °C
- Relative Humidity (RH): 10–90%
- Nozzle Temperature (T_{nozzle}): 200–230 °C
- Infill Density (ρ_{infill}): 80–100%
- Layer Thickness (h_{layer}): 0.15–0.25 mm
- Printing Speed (V_{print}): 50–70 mm/s

Within these bounds, temporal trends and abrupt excursions were programmatically introduced to emulate disturbance scenarios. Note that parameters such as layer thickness and infill density, typically static in standard G-code, were allowed to vary dynamically in this simulation to represent inter-layer adaptive slicing scenarios and to maximize the entropy of the data stream for stress-testing the ingestion pipeline.

A total of 100,000 synthetic records were generated at fixed 2 s intervals. The data were exported in CSV format for compatibility purposes. Table 3 presents an illustrative excerpt of the generated telemetry stream.

The resulting dataset functions as a digital testbed for evaluating the computational throughput, responsiveness, and anomaly detection capabilities. Future work will incorporate experimentally acquired signals and transitioning to hybrid cyber–physical datasets, strengthening the link between computational indicators and physical process behavior.

Table 3. Excerpt of the generated telemetry stream showing heuristic signal coupling. Values display high variance to verify the system’s dynamic range handling. By Record 3, the system emulates a ‘cold extrusion’ risk by correlating low Nozzle Temp. (202.04 °C) with reduced Printing Speed (51.80 mm/s).

Feature	Rec. 1	Rec. 2	Rec. 3	Rec. 4	Rec. 5
Ambient Temp. (°C)	25.83	31.51	29.57	30.69	31.74
Relative Humidity (%)	72.67	89.72	85.64	53.19	24.47
Nozzle Temp. (°C)	221.55	212.13	202.04	211.68	226.75
Infill Density (%)	92.56	88.56	95.57	93.39	95.49
Layer Thickness (mm)	0.172	0.227	0.157	0.243	0.159
Printing Speed (mm/s)	62.21	60.66	51.80	52.44	68.26

4.4. Performance Evaluation and Computational Benchmarking

The dataset was used to benchmark all system layers under consistent initial conditions, enabling a quantitative assessment of the throughput, latency, and resource efficiency. Each architectural component was evaluated as follows.

4.4.1. Data Ingestion Layer

Simulated IoT data pipelines were established using Apache Kafka and RabbitMQ version 4.0.7 to enable high-frequency message streaming [45,46]. Although this enterprise-grade stack is more robust than necessary for the current dataset volume, it was chosen to showcase architectural readiness for extensive parallel scaling across multi-machine server farms. The ingestion modules utilized were as follows:

- Batched message buffering to minimize transmission overhead
- Adaptive recovery mechanisms for load balancing and fault tolerance
- Asynchronous queuing to sustain stable throughput under varying data rates

Consequently, the Data Ingestion Layer is configured to subscribe to a multi-channel telemetry stream. The input vector V_t is defined to include the core rheological, kinematic, and environmental parameters:

$$V_t = \{T_{amb}, RH, T_{nozzle}, \rho_{infill}, h_{layer}, V_{print}\}$$

The vector captures the critical interplay between thermal conditions (T_{nozzle}, T_{amb}), extrusion dynamics ($V_{print}, h_{layer}, \rho_{infill}$), and external disturbances (RH). These variables are sufficient to diagnose the majority of common failure modes (e.g., thermal instability, layer shift, moisture degradation). The Kafka producers serialize these values into JSON payloads at a frequency of 0.5 Hz.

4.4.2. Processing and Analytics Layer

Data analysis combined real-time and batch-processing paradigms. In particular:

- Apache Spark Streaming was used for live anomaly detection
- Apache Hadoop enabled long-term data aggregation and offline diagnostics [47–49]

This dual-engine approach demonstrates the architecture’s capability to handle both immediate alert generation and historical trend analysis, a prerequisite for digital twin implementation. The deterministic data ensured that each run produced comparable latency and throughput metrics, allowing for reliable cross-validation of computational efficiency.

4.4.3. Application Layer

The application interface translates analytical outcomes into operator alerts and dashboard visualizations. Two configurations were evaluated.

- A monolithic structure (Algorithm A1)
- A modular microservice deployment using Docker containers (Algorithms A2 and A3) [50]

The distributed implementation achieved higher scalability and robustness, successfully issuing early warning notifications when the RH exceeded the 50% threshold, simulating the onset of PLA degradation [9,51].

4.5. Key Performance Indicators

The framework performance was quantified using computational Key Performance Indicators (KPIs) relevant to real-time monitoring environments.

- Throughput (messages/s): data-handling efficiency
- End-to-End Processing Time (s): latency from ingestion to alert output
- CPU and Memory Usage (%): system resource utilization
- Data Volume (MB): scalability and storage demand
- Average Latency per Batch (s): responsiveness in near-real-time detection
- Processing Time per Message (s): per-sample computational efficiency

These indicators provide a quantitative baseline for evaluating the process monitoring performance of FFF. The seed-controlled data generation approach ensures that all analyses are reproducible, allowing for transparent comparisons across configurations and future integration with real-world process data.

Subsequent research will extend this digital verification to physical experimentation, linking digital anomaly patterns to observable defects such as warping, delamination, and surface irregularities.

5. Numerical Results

In this section, we quantitatively evaluate the proposed early-warning monitoring framework using a simulated FFF dataset. Our analysis focuses on the computational performance of the system across its three architectural layers—data ingestion, data processing, and application—under conditions that simulate continuous manufacturing environments.

The obtained numerical results showed the capability of the framework to process telemetry data faster than the thermal characteristic time of the polymer extrusion. This ensures that alerts are triggered before irreversible defects, such as crystallization and hydrolysis, occur in the final part. Hence, the framework provides scalable, data-driven environmental and process supervision, offering early warning insights rather than direct control interventions. At this stage, physical manufacturing indicators such as dimensional accuracy or tensile strength are excluded and will be introduced in the next phase through experimental validation using dedicated sensor hardware.

In the following subsections, we detail the performance of each architectural layer, focusing on the throughput, latency, and resource utilization. The metrics summarized in Tables 4–8 collectively affirm the scalability, responsiveness, and reliability of the proposed simulation-based monitoring system, which is designed for integration into data-intensive and sustainable manufacturing infrastructure.

5.1. Evaluation of the Data Ingestion Layer

The data ingestion layer was benchmarked to assess its capacity for high-frequency data capture by replicating the IoT-enabled AM environments. Two industrial messaging

systems, RabbitMQ and Apache Kafka, were evaluated under identical operating conditions, each managing 100,000 simulated sensor records containing T, RH, and nozzle parameter data.

Reliable ingestion is essential for real-time monitoring pipelines to ensure uninterrupted data flow toward the analytics and anomaly detection modules. The results summarized in Table 4 show significant differences in throughput and latency performance between the two systems.

Table 4. Performance metrics comparison between RabbitMQ (RMQ) and Apache Kafka under identical simulation conditions.

Metric	RMQ Prod.	RMQ Cons.	Kafka Prod.	Kafka Cons.
Total Time (s)	29.87	34.43	7.79	16.25
Throughput (msg/s)	3347.84	2904.45	12,830.58	6155.71
Avg. Latency (s)	2.1076	–	0.0769	–
Avg. Proc. Time (s)	–	0.000001	–	0.0000004

In our study, Apache Kafka showed optimal scalability and responsiveness, achieving a producer throughput of 12,830.58 msg/s. This numerical result is nearly four times greater than that of RabbitMQ, along with an average latency of 0.0769 s. In the context of FFF, where a print head moving at 60 mm/s covers almost 5 mm in 0.1 s, minimizing latency is essential. Kafka’s sub-second response time ensures that environmental spikes are detected almost instantaneously, whereas RabbitMQ’s 2-s delay creates a blind spot where transient environmental spikes could be missed or where the machine executes numerous instructions without supervision. Although the experimental sampling rate was conservatively set to 0.5 Hz, the near-zero latency of the ingestion layer is critical to ensure that the monitoring architecture itself introduces no additional lag, keeping the total system reaction time strictly dependent on the physical sensor sampling rather than on computational overhead.

This high performance underscores the suitability of Kafka for near-real-time process supervision, where continuous monitoring of temperature and RH variations is essential to prevent deviations that could affect part integrity. Figure 4 graphically compares the ingestion performance, highlighting the stability and efficiency of the proposed architecture for IoT-enabled proactive process management.

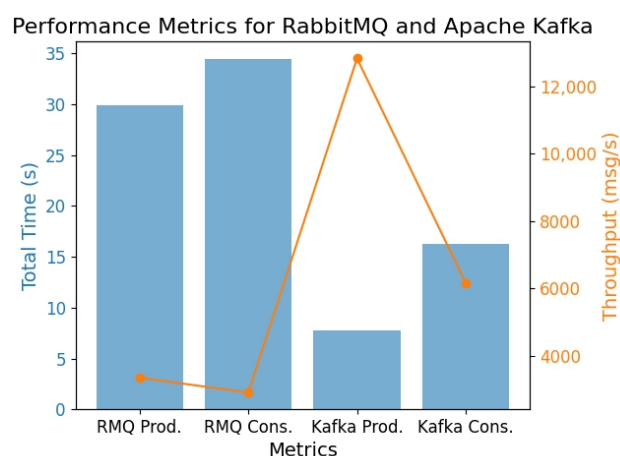


Figure 4. RabbitMQ vs. Apache Kafka in terms of total time and throughput (our elaboration).

These results confirm that Kafka’s distributed event streaming model offers the responsiveness and fault tolerance necessary for industrial-grade AM monitoring pipelines, aligning with the principles of smart and sustainable manufacturing.

5.2. Semantic Correlation of System Alerts

While the architectural scope of this study excludes the physical simulation of microstructural damage, the monitoring framework provides significant semantic value by mapping detected signal anomalies to specific defect risks. By leveraging the heuristic logic defined in Section 4.1, Table 5 presents the classification schema applied to the alerts generated during the tests.

Table 5. Logic Mapping between Telemetry Triggers, System Alerts, and Associated Quality Risks.

Trigger Condition	System Alert Output	Associated Quality Risk
$RH > 50\%$ (sustained)	WARN_HUMIDITY_CRIT	Hydrolysis/Porosity Risk
$T_{nozzle} < T_{min}$	ERR_THERMAL_UNDERRUN	Cold Extrusion/Clogging
$\Delta T_{amb} / \Delta t > Threshold$	WARN_THERMAL_SHOCK	Warping (Thermal Contraction)
$Flow \approx 0$	CRIT_FLOW_STALL	Deposition Failure (Filament Break)

This classification demonstrates that the architecture is capable of transforming raw sensor signals into actionable process knowledge. Even without in situ metrology, the system provides operators with an explicit indication of the type of defect likely to occur (e.g., distinguishing between a thermal drift and a flow interruption), thus demonstrating the diagnostic utility of the proposed heuristics.

5.3. Evaluation of the Data Processing Layer

The data processing layer was analyzed by comparing two complementary paradigms: batch analytics (Apache Hadoop) and real-time stream analytics (Apache Spark). This evaluation enables the assessment of the trade-offs between historical data aggregation and continuous process supervision.

Table 6. Performance metrics of Hadoop batch analytics for simulated FFF process data.

Metric	Value
Total Records Processed	100,000
Processing Time (s)	4.640
Throughput (records/s)	21,550.46
CPU Usage (Before/After)	43.3%/47.0%
Memory Usage (Before/After)	64.9%/64.3%
Avg. Temperature (C)	27.49
Avg. Relative Humidity (%)	49.93

Hadoop efficiently processed all 100,000 simulated records within 4.64 s, achieving a throughput of over 21,500 records/s. Although highly efficient for retrospective analytics, its batch-oriented structure introduces inherent delays, thereby limiting its suitability for real-time anomaly detection and immediate operator feedback. Conversely, Spark Streaming showed the ability to detect process anomalies in near real time while maintaining computational efficiency. The processing metrics of the two consecutive micro-batches are presented in Table 7.

Table 7. Performance metrics for Apache Spark Streaming across consecutive micro-batches.

Metric	Batch ID: 0	Batch ID: 1	Summary
Record Count	3000	1001	4001
Processing Time (s)	1.897	1.247	3.145
Throughput (records/s)	1581.08	802.47	-
Latency (s)	1.897	1.247	1.572 (avg)
CPU Usage (Before/After)	0.0%/30.1%	32.0%/20.5%	-
Memory Usage (Before/After)	67.2%/67.1%	67.1%/67.1%	-
Avg. Temperature (C)	27.59	27.73	27.66
Avg. Relative Humidity (%)	49.15	49.90	49.53

With reference to Spark Streaming, this system achieved an average latency of 1.57 s, confirming its adequacy for the continuous supervision of dynamic parameters such as RH and nozzle temperature. By observing that humidity-induced degradation in PLA is a diffusive process that evolves over minutes rather than milliseconds, a 1.5-s detection window provides an ample safety margin for operator intervention in order to prevent defects in the final parts. This responsiveness ensures that deviations beyond the 50% RH threshold can be detected early enough to trigger preventive actions, thereby reducing the likelihood of print defects.

The combination of batch and stream analytics thus provides complementary capabilities: Hadoop supports large-scale historical evaluation and trend identification, whereas Spark enables near real-time decision support for adaptive control in AM processes.

Figure 5 compares the computational performance in terms of processing time for Apache Hadoop and Apache Spark Streaming under identical synthetic workloads. Spark Streaming exhibited a lower processing time (3.145 s) than Hadoop (4.640 s), confirming its suitability for near-real-time analytics tasks, where latency is critical. Moreover, while Hadoop requires the accumulation of the entire dataset (100,000 records) before execution, resulting in a high initial latency barrier, Spark Streaming processes incoming micro-batches (e.g., 3000 records) in just 1.897 s. This demonstrates that Spark provides actionable insights immediately during the process, whereas Hadoop is superior only for post-process aggregated throughput.

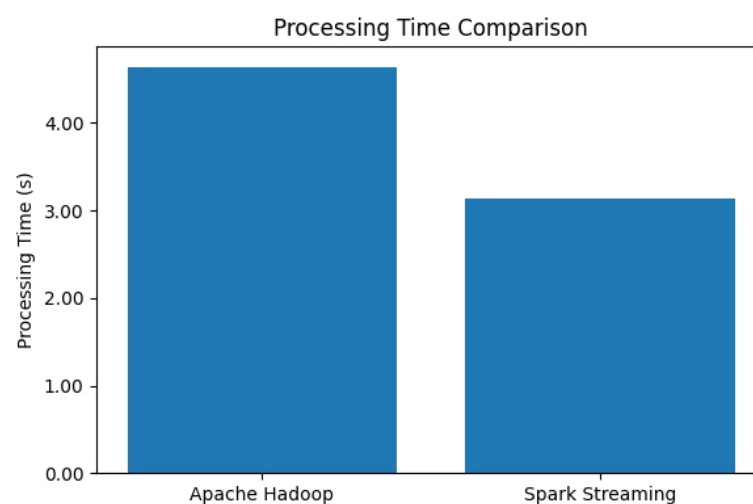


Figure 5. Processing latency comparison between batch (Hadoop) and stream (Spark) analytics. Spark Streaming significantly reduces reaction time by processing micro-batches, whereas Hadoop's batch initialization overhead limits its utility for immediate anomaly detection.

However, computational efficiency must be balanced with resource consumption. As illustrated in Figures 6 and 7, Hadoop demonstrates a specific resource utilization profile reflecting its optimization for high-throughput batch processing. Specifically, Figure 6 compares the CPU load before and during the task, showing the distinct operational demands of the two frameworks. Similarly, memory utilization (Figure 7) indicates that both frameworks are memory-intensive, with Spark requiring substantial allocation to support its in-memory processing architecture.

These results highlight the complementary roles of batch and streaming frameworks within the proposed architecture: Spark Streaming enables the rapid anomaly detection required for the early-warning framework (as further visualized in the anomaly detection case study), while Hadoop supports efficient retrospective analysis and process auditing.

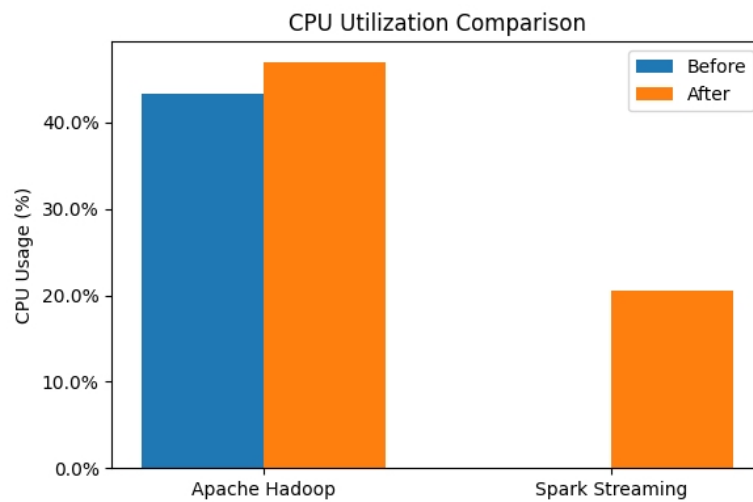


Figure 6. CPU resource utilization profiles. The difference in processor load highlights the distinct operational demands of Hadoop’s intensive batch aggregation versus the continuous, real-time computational flow of Spark Streaming.




Figure 7. Memory consumption analysis. Both frameworks exhibit high memory demand, with Spark’s in-memory computing architecture requiring substantial allocation to maintain data state for rapid, low-latency stream processing.

5.4. Evaluation of the Application Layer

The application layer governs the orchestration and alerting functions of the framework. Two implementation strategies were evaluated under identical simulation loads:

a monolithic application and a microservice-based architecture deployed using Docker containers. Each configuration processed 100,000 records divided into ten batches. The results are presented in Table 8.

Table 8. Trade-off analysis: Monolithic raw speed vs. Microservice resilience.

Metric	Monolithic	Microservice
Total Batches Processed	10	10
Total Records Processed	100,000	100,000
Avg. Response Time (s)	1.05	1.98
Avg. Throughput (records/s)	9705.54	5040.52
Total Failed Requests	0	0

Although the monolithic configuration achieved higher raw throughput and shorter response times, the microservice-based implementation is considered the better choice for industrial scaling. Despite a slight increase in latency (1.98 s), which remains well within the safety margins for environmental monitoring, the microservice architecture effectively eliminates single points of failure (SPOF). In a real-world scenario involving numerous printers, a crash in one analysis module would render the entire farm blind in a monolithic system. In contrast, microservices isolate faults, allowing the rest of the production line to remain under supervision. Consequently, microservice-based implementation offers greater modularity, resilience, and ease of integration, which are essential for enabling CM and Industry 4.0 ecosystems. This distributed architecture supports fault tolerance, independent service scaling, and interoperability with external modules such as digital twins and AI-based controllers.

Overall, the computational evaluation confirmed that the proposed architecture effectively handled large-scale, high-frequency process data while maintaining a responsiveness that was suitable for early warning supervision. The results establish a verified computational foundation for subsequent experimental integration, where simulated anomalies are correlated with actual process outcomes such as warping, interlayer delamination, or mechanical degradation.

5.5. Note on Detection Metrics

It is relevant to address the absence of standard classification metrics (e.g., Sensitivity, False Positive Rate, F1-Score) in this evaluation. These metrics require a labeled ‘ground truth’ dataset that correlates sensor anomalies with confirmed physical defects on printed parts. Since this study focuses on the verification of the computational architecture and data pipeline using simulated process variables, such statistical metrics would be artificial.

In the proposed deterministic framework, detection performance is strictly a function of system reliability: if the architecture successfully transmits the data packet without loss (as indicated by the Latency and Jitter results), the rule-based logic identifies the deviation by definition. Therefore, the architectural benchmarks presented above (Latency < 2 s, Throughput stability) serve as the primary indicators of operational readiness for this phase of research.

6. Discussion

Our experimental findings demonstrate the technical feasibility of a simulation-driven computational framework for early-warning and anticipatory supervision in FFF. Importantly, the findings should be interpreted as a verification of the computational backbone rather than as evidence of fully validated material science predictions. The contribution of this work lies in establishing the data-handling, analytics, and decision-latency char-

acteristics required for near real-time (low-latency) supervision, serving as the necessary architectural precursor to physical deployment.

6.1. Analytics Performance and Semantic Value

Within this scope, the framework demonstrated a dual capability: high-speed data ingestion and semantic interpretation of anomalies. As detailed in the Semantic Correlation analysis (Table 5), the system goes beyond generic error logging by effectively mapping raw telemetry violations to specific quality risks (e.g., linking humidity drifts to hydrolysis risk). Quantitatively, the architecture supports this logic with high throughput: Apache Kafka processes over 12,830 messages/s with a latency below 0.1 s (0.0769 s). From a supervisory perspective, this confirms that the bottleneck for closed-loop control is not the software stack, but the physical process dynamics, confirming the viability of general-purpose IoT protocols for industrial monitoring.

The comparative analysis between Hadoop and Spark Streaming further clarifies the architectural roles. While Hadoop (approx. 21,550 records/s) serves retrospective auditing, Spark Streaming (1.57 s latency) provides the responsiveness required for the “Intra-layer Intervention Window.” While this latency does not guarantee the prevention of instantaneous failures, it provides sufficient actionable lead time for moisture-sensitive materials, where degradation kinetics unfold over minutes rather than milliseconds.

6.2. Operational Readiness: Latency vs. Physical Dynamics

To justify the system’s early-warning potential, we quantify the operational margin by comparing computational latency against the physical inertia of the FFF process. A critical metric is the thermal time constant of the liquefier assembly, which acts as a physical low-pass filter against rapid temperature changes. As modeled in foundational studies on FFF dynamics [52], the heating system exhibits a characteristic response time that significantly exceeds the millisecond scale of electronic data transmission.

Comparing this physical inertia (typically in the order of seconds) with our measured system latency ($t_{sys} \approx 1.98$ s) confirms that the monitoring architecture operates within a safe temporal margin. The response ratio t_{sys}/τ remains well below unity, theoretically allowing the detection of thermal drift before the nozzle temperature physically deviates enough to cause rheological instabilities.

6.3. Architectural Resilience and Circular Economy Implications

At the application layer, the trade-off between monolithic and microservice architectures has significant implications for autonomous manufacturing. Although the monolithic configuration achieved marginally lower latency (1.05 s vs. 1.98 s), the microservice architecture offers the fault isolation required for “lights-out” operations. A single service failure in a monolithic block risks halting the entire supervision line, whereas the distributed approach ensures system survivability across diverse operational segments.

Furthermore, the proposed framework aligns with Circular Economy principles by prioritizing resource efficiency and waste minimization. By monitoring environmental parameters such as Ambient Temperature and RH in real time, the platform provides actionable information to prevent irreversible material degradation. This shift from post-process rejection to in-process risk mitigation directly contributes to a more sustainable manufacturing paradigm, reducing polymer waste and optimizing energy usage.

6.4. Limitations and Scope of Validity

It is essential to acknowledge the limitations of this architectural study. First, the verification relies on synthetic data injection rather than physical experimentation. Although the anomaly patterns were derived from heuristic logic (Section 4.1) and mapped to semantic

risks (Section 5.2), the framework does not yet incorporate a coupled thermo-mechanical model. As a result, the term “predictive” is strictly limited to the data-driven anticipation of risk states, not to the first-principles prediction of final part geometry.

Future work will focus on experimentally validating these architectural findings. This will involve correlating the generated alerts with the physical measurements of dimensional accuracy and interlayer strength on a real FFF testbed. Establishing this ground truth is the critical next step to translate the current “Operational Readiness” into a fully calibrated quality assurance system.

7. Conclusions

This study presented a process-oriented computational framework for monitoring and early-warning supervision in FFF, designed to bridge the gap between high-frequency sensor data and actionable process insights. By prioritizing a microservice-based architecture, the study focused on verifying the computational feasibility of soft real-time supervision under FFF-relevant data rates. Simulation-based results confirmed that high-throughput data ingestion and distributed stream analytics can be effectively integrated to support the low-latency detection of environmentally induced process instabilities.

Benchmarking results indicate that the combined use of Apache Kafka and Spark Streaming enables sub-second to near-real-time responsiveness. This performance ensures that the system can operate well within the thermal time constants of the extrusion process, preventing data transmission delays from bottlenecking the detection logic. These findings demonstrate that the proposed architecture is computationally capable of supporting intervention specifically within the intra-layer deposition window, before adverse environmental conditions escalate into irreversible process failure.

The adoption of a simulation-first methodology serves as a strategic de-risking step. Unlike physical testbeds, which are constrained by material costs and hardware safety, the virtual environment allowed for aggressive stress-testing of the data pipeline under extreme load conditions. While this approach enables a systematic evaluation of throughput and resilience, the conclusions remain confined to the digital domain. The absence of hardware-in-the-loop testing means the current framework provides data-driven anticipation of instability rather than a physically validated prediction of part quality. By shifting the paradigm from post-process inspection to in-process risk mitigation, the system aligns with Circular Economy principles and proactively minimizes material and energy waste.

Future work will focus on a proof-of-concept implementation on an actual FFF system, including integration with physical sensors, end-to-end latency assessment, and the correlation of early-warning indicators with measurable defects such as dimensional inaccuracies, warping, and interlayer adhesion loss. Additional studies will investigate the influence of controlled pre-print RH exposure on filament condition to quantitatively establish the link between environmental history and observable artifacts.

Overall, this work provides a verified computational foundation for early-warning supervision in extrusion-based AM. By delineating verified capabilities alongside current limitations, it provides a structured progression from simulation-based architectural validation toward fully validated, cyber-physical predictive monitoring in FFF systems.

Author Contributions: Conceptualization, M.P.; methodology, M.P. and A.P.; validation, M.P. and A.P.; resources, G.P.; writing—original draft preparation, A.P.; writing—review and editing, M.P.; supervision, M.P. and G.P.; project administration, G.P.; funding acquisition, M.P. and G.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The used data can be shared upon request.

Acknowledgments: The authors acknowledge the administrative and technical support provided by Advantech S.r.l. Industry, Lecce, Italy.

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

Abbreviations

The following abbreviations are used in this manuscript:

AM	Additive Manufacturing
API	Application Programming Interface
CAD	Computer-Aided Design
CM	Cloud Manufacturing
FDM	Fused Deposition Modeling
FFF	Fused Filament Fabrication
IoT	Internet of Things
KPI	Key Performance Indicator
PLA	Polylactic Acid
RH	Relative Humidity
SPOF	Single Point of Failure
STL	Stereolithography

Appendix A. Experimental Results

In these sections, we report the obtained results in terms of the average parameter values and warning messages for both the monolithic and microservice architectures.

Appendix A.1. Monolithic App Results

Table A1 lists the processed batch data for the monolithic application. Note that the system flagged specific batches where environmental parameters deviated from the safety thresholds.

Table A1. Average values and warning messages for the monolithic app.

Batch ID	Avg. T (C)	Avg. RH (%)	Avg. Nozzle T (C)	Avg. Infill (%)	Avg. Layer (mm)	Avg. Speed (mm/s)	System Message
1	27.54	50.30	215.10	95.2	0.200	60.01	WARNING: High humidity might affect mechanical properties.
2	27.43	49.92	214.85	97.2	0.201	59.95	Optimal range.
3	27.51	49.80	215.08	93.0	0.200	60.05	Optimal range.
4	27.40	49.95	215.15	98.5	0.200	60.07	Optimal range.
5	27.48	49.78	215.00	91.0	0.199	59.88	Optimal range.
6	27.47	49.96	215.02	96.7	0.200	60.00	Optimal range.
7	27.46	50.25	214.92	94.0	0.200	59.97	WARNING: High humidity might affect mechanical properties.
8	27.50	49.93	214.96	97.0	0.199	60.11	Optimal range.
9	27.55	49.73	214.90	93.2	0.200	59.94	Optimal range.
10	27.53	49.72	215.01	90.1	0.200	59.90	Optimal range.

Appendix A.2. Microservice Architecture Results

Table A2 presents the results for the microservice-based implementation. Due to different logging protocols, the “None” status in this table corresponds to the “Optimal range” condition in the monolithic system.

Table A2. Average values and warnings for the microservice app.

Batch ID	Avg. T (C)	Avg. RH (%)	Avg. Nozzle T (C)	Avg. Infill (%)	Avg. Layer (mm)	Avg. Speed (mm/s)	System Message
1	27.52	50.28	215.05	95.0	0.200	60.00	Relative humidity exceeded 50% in Batch 1
2	27.46	49.88	214.92	97.4	0.200	59.96	None
3	27.48	49.85	215.03	92.7	0.200	60.01	None
4	27.43	49.98	215.10	98.3	0.200	60.05	None
5	27.51	49.75	215.06	90.9	0.199	59.89	None
6	27.50	49.91	214.97	96.6	0.200	59.97	None
7	27.44	50.24	214.94	94.2	0.200	59.98	Relative humidity exceeded 50% in Batch 7
8	27.49	49.95	214.92	97.3	0.199	60.08	None
9	27.52	49.70	214.88	93.0	0.200	59.93	None
10	27.50	49.73	214.97	90.0	0.200	59.91	None

Appendix A.3. Functional Evidence of the Alerting Capabilities

An analysis of the processed data revealed that both the monolithic and microservice architectures maintained operational averages for key FFF parameters, such as ambient/nozzle temperature, infill density, layer thickness, and printing speed. Although slight numerical differences are evident between the two tables owing to the distinct processing pipelines, the results remain reasonably close, reinforcing that both architectures provide reliable interpretations of the machine and process data.

Both architectures successfully flagged critical deviations in the environmental conditions. The monolithic system triggered high humidity warnings for Batches 1 and 7. Similarly, the microservice architecture identified and flagged the same batches (1 and 7) for exceeding the 50%-RH threshold. The accurate detection of these humidity spikes, which are known to cause filament degradation, serves as functional verification of the computational methodology. This confirms the capacity of the system to proactively detect quality risks, thereby ensuring similar algorithmic performance across both architectures.

Appendix B. Algorithmic Implementation

In this appendix, we detail the algorithmic implementation of the Application Layer, which serves as the core decision-support engine of the proposed framework. The pseudocodes provided below illustrate the operational logic responsible for ingesting aggregated telemetry, evaluating critical parameters against safety thresholds (e.g., RH_{limit}), and triggering real-time early-warning alerts.

To clarify the architectural distinctions discussed in Section 5, we present the logic for both the centralized (monolithic) and distributed (microservice-based) implementations in our framework. Following Algorithm A1 outlines the unified workflow for the monolithic system, where data processing and alerting occur within a single sequential execution cycle. Algorithms A2 and A3 implement the decoupled logic of the microservice architecture, highlighting the separation of concerns between data processing and system performance monitoring.

Algorithm A1 Monolithic Application Logic

```

1: Initialize Spark session and performance metrics containers
2: Load telemetry data source (IoT_3D_printing_data.csv)
3: Define batch processing size and safety thresholds ( $RH_{limit} = 50\%$ )
4: for each batch  $b$  in data stream do
5:   Compute statistical averages for batch  $b$ 
6:   if  $RH_{avg} > RH_{limit}$  then
7:     Trigger Warning: "High humidity detected"
8:   else
9:     Log Status: "Optimal range"
10:  end if
11:  Append results to output log
12: end for
13: Export final performance metrics

```

Algorithm A2 Data Processing Microservice (Flask)

```

1: Initialize Flask application context
2: Set global thresholds ( $T_{limit}, RH_{limit}$ )
3: function PROCESSBATCH(incoming_json)
4:    $stats \leftarrow$  Compute averages from incoming_json
5:    $warnings \leftarrow$  Evaluate  $stats$  against thresholds
6:   Forward metrics to Monitoring Service API
7:   return JSON response with  $warnings$ 
8: end function
9: Run server on port 5002

```

Algorithm A3 Microservice App Metrics Collection

```

1: Define target service URLs (Ingestion, Processing)
2: function COLLECTMETRICS
3:   try
4:      $metrics \leftarrow$  Poll APIs for CPU, RAM, Latency
5:     Store  $metrics$  in time-series database
6:   catch ConnectionError
7:     Log error and retry with exponential backoff
8: end function
9: Schedule COLLECTMETRICS every  $N$  seconds

```

References

- Hull, C.W. Apparatus for Production of Three-Dimensional Objects by Stereolithography. U.S. Patent 4,575,330A, 11 March 1986.
- Zgodavová, K.; Lengyelová, K.; Bober, P.; Eguren, J.A.; Moreno, A. 3D printing optimization for environmental sustainability: Experimenting with materials of protective face shield frames. *Materials* **2021**, *14*, 6595. [[CrossRef](#)]
- ISO/ASTM 52900:2021; Additive manufacturing—General principles—Fundamentals and vocabulary. International Organization for Standardization: Geneva, Switzerland, 2021.
- Kawalkar, R.; Dubey, H.K.; Lokhande, S.P. A review for advancements in standardization for additive manufacturing. *Mater. Today Proc.* **2022**, *50*, 1983–1990. [[CrossRef](#)]
- Desai, A.A.; Patil, S.D.; Yadav, P.H.; Kekare, A. Design, Analysis, and Development of Additive Manufacturing by using FDM Technique with Dual Nozzle Assembly. *Mater. Today Proc.* **2023**, *77*, 619–626. [[CrossRef](#)]
- Hussain, S.; Kallaste, A.; Vaimann, T. Recent Trends in Additive Manufacturing and Topology Optimization of Reluctance Machines. *Energies* **2023**, *16*, 3840. [[CrossRef](#)]
- Butt, J.; Bhaskar, R.; Mohaghegh, V. Analysing the effects of layer heights and line widths on FFF-printed thermoplastics. *Int. J. Adv. Manuf. Technol.* **2022**, *121*, 7383–7411. [[CrossRef](#)]
- Muthe, L.P.; Pickering, K.; Gauss, C. A review of 3D/4D printing of poly-lactic acid composites with bio-derived reinforcements. *Compos. Part C Open Access* **2022**, *8*, 100271. [[CrossRef](#)]

9. Lendvai, L.; Fekete, I.; Jakab, S.K.; Szarka, G.; Verebélyi, K.; Iván, B. Influence of environmental humidity during filament storage on the structural and mechanical properties of material extrusion 3D-printed poly (lactic acid) parts. *Results Eng.* **2024**, *24*, 103013. [[CrossRef](#)]
10. Ghorbani, J.; Koirala, P.; Shen, Y.L.; Tehrani, M. Eliminating voids and reducing mechanical anisotropy in fused filament fabrication parts by adjusting the filament extrusion rate. *J. Manuf. Processes* **2022**, *80*, 651–658. [[CrossRef](#)]
11. Kattinger, J.; Kornely, M.; Ehrler, J.; Bonten, C.; Kreutzbruck, M. Analysis of melting and flow in the hot-end of a material extrusion 3D printer using X-ray computed tomography. *Addit. Manuf.* **2023**, *76*, 103762. [[CrossRef](#)]
12. Zhai, C.; Wang, J.; Tu, Y.P.; Chang, G.; Ren, X.; Ding, C. Robust optimization of 3D printing process parameters considering process stability and production efficiency. *Addit. Manuf.* **2023**, *71*, 103588. [[CrossRef](#)]
13. Velghe, I.; Buffel, B.; Vandeginste, V.; Thielemans, W.; Desplentere, F. Review on the Degradation of Poly (lactic acid) during Melt Processing. *Polymers* **2023**, *15*, 2047. [[CrossRef](#)]
14. Fang, L.; Yan, Y.; Agarwal, O.; Yao, S.; Seppala, J.E.; Kang, S.H. Effects of environmental temperature and humidity on the geometry and strength of polycarbonate specimens prepared by fused filament fabrication. *Materials* **2020**, *13*, 4414. [[CrossRef](#)]
15. Banjo, A.D.; Agrawal, V.; Auad, M.L.; Celestine, A.D.N. Moisture-induced changes in the mechanical behavior of 3D printed polymers. *Compos. Part C Open Access* **2022**, *7*, 100243. [[CrossRef](#)]
16. Laumann, D.; Spiehl, D.; Dörsam, E. Influence of printing material moisture on part adhesion in fused filament fabrication 3D printing. *J. Adhes.* **2024**, *100*, 380–394. [[CrossRef](#)]
17. Sola, A.; Rosa, R.; Ferrari, A.M. Environmental Impact of Fused Filament Fabrication: What Is Known from Life Cycle Assessment? *Polymers* **2024**, *16*, 1986. [[CrossRef](#)]
18. Hsueh, M.H.; Lai, C.J.; Wang, S.H.; Zeng, Y.S.; Hsieh, C.H.; Pan, C.Y.; Huang, W.C. Effect of printing parameters on the thermal and mechanical properties of 3d-printed pla and petg, using fused deposition modeling. *Polymers* **2021**, *13*, 1758. [[CrossRef](#)] [[PubMed](#)]
19. Aniskevich, A.; Bulderberga, O.; Stankevics, L. Moisture sorption and degradation of polymer filaments used in 3D printing. *Polymers* **2023**, *15*, 2600. [[CrossRef](#)]
20. Agrawal, S.; Oza, P.; Kakkar, R.; Tanwar, S.; Jetani, V.; Undhad, J.; Singh, A. Analysis and recommendation system-based on PRISMA checklist to write systematic review. *Assess. Writ.* **2024**, *61*, 100866. [[CrossRef](#)]
21. Mazzarisi, M.; Angelastro, A.; Campanelli, S.L.; Errico, V.; Posa, P.; Fusco, A.; Colucci, T.; Edwards, A.J.; Corigliano, S. Monitoring of Directed Energy Deposition Laser Beam of Nickel-Based Superalloy via High-Speed Mid-Wave Infrared Coaxial Camera. *J. Manuf. Mater. Process.* **2024**, *8*, 294. [[CrossRef](#)]
22. Borghetti, M.; Cantù, E.; Sardini, E.; Serpelloni, M. Future sensors for smart objects by printing technologies in Industry 4.0 scenario. *Energies* **2020**, *13*, 5916. [[CrossRef](#)]
23. Fedullo, T.; Morato, A.; Peserico, G.; Trevisan, L.; Tramarin, F.; Vitturi, S.; Rovati, L. An IoT measurement system based on LoRaWAN for additive manufacturing. *Sensors* **2022**, *22*, 5466. [[CrossRef](#)] [[PubMed](#)]
24. Horr, A.M.; Drexler, H. Real-Time Models for Manufacturing Processes: How to Build Predictive Reduced Models. *Processes* **2025**, *13*, 252. [[CrossRef](#)]
25. Mao, Y.; Lin, H.; Yu, C.X.; Frye, R.; Beckett, D.; Anderson, K.; Jacquemetton, L.; Carter, F.; Gao, Z.; Liao, W.k.; et al. A deep learning framework for layer-wise porosity prediction in metal powder bed fusion using thermal signatures. *J. Intell. Manuf.* **2023**, *34*, 315–329. [[CrossRef](#)]
26. Knaak, C.; Masseling, L.; Duong, E.; Abels, P.; Gillner, A. Improving build quality in laser powder bed fusion using high dynamic range imaging and model-based reinforcement learning. *IEEE Access* **2021**, *9*, 55214–55231. [[CrossRef](#)]
27. Takalo-Mattila, J.; Heiskanen, M.; Kyllönen, V.; Määttä, L.; Bogdanoff, A. Explainable steel quality prediction system based on gradient boosting decision trees. *IEEE Access* **2022**, *10*, 68099–68110. [[CrossRef](#)]
28. Aniolis, J.; Dudzik, G.; Abramski, K.M. Automated non-destructive 3D printing filament material properties monitoring based on electric permittivity, longitudinal encoding and diameter multi-axes real-time measurements. *Addit. Manuf.* **2024**, *84*, 104103. [[CrossRef](#)]
29. Devito, F.; Mazzarisi, M.; Dassisti, M.; Lavecchia, F. Causal technological model for predicting void fraction and energy consumption in material extrusion process of polylactic acid. *J. Manuf. Processes* **2024**, *129*, 187–201. [[CrossRef](#)]
30. Yang, J.; Liu, Y. Hybrid modelling method for the prediction and experimental validation of 3D printing resource consumption. *J. Manuf. Processes* **2023**, *101*, 1275–1300. [[CrossRef](#)]
31. Li, X.; Zhang, M.; Zhou, M.; Wang, J.; Zhu, W.; Wu, C.; Zhang, X. Qualify assessment for extrusion-based additive manufacturing with 3D scan and machine learning. *J. Manuf. Processes* **2023**, *90*, 274–285. [[CrossRef](#)]
32. Scime, L.; Siddel, D.; Baird, S.; Paquit, V. Layer-wise anomaly detection and classification for powder bed additive manufacturing processes: A machine-agnostic algorithm for real-time pixel-wise semantic segmentation. *Addit. Manuf.* **2020**, *36*, 101453. [[CrossRef](#)]

33. Breseghello, L.; Naboni, R. Toolpath-based design for 3D concrete printing of carbon-efficient architectural structures. *Addit. Manuf.* **2022**, *56*, 102872. [[CrossRef](#)]
34. de Kergariou, C.; Kim, B.C.; Perriman, A.; Le Duigou, A.; Guessasma, S.; Scarpa, F. Design of 3D and 4D printed continuous fibre composites via an evolutionary algorithm and voxel-based finite elements: Application to natural fibre hygromorphs. *Addit. Manuf.* **2022**, *59*, 103144. [[CrossRef](#)]
35. Engel, S.; Hegger, J.; Classen, M. Multimodal automated fabrication with concrete: Case study and structural performance of ribbed CFRP-reinforced concrete ceiling. *Addit. Manuf.* **2025**, *101*, 104689. [[CrossRef](#)]
36. Wu, J.; Guo, P.; Zhang, H.; Mai, T.; Zhang, K.; Li, A.; Lyu, Y.; Zhang, H.; Yang, D. Material extrusion additive manufacturing of recycled discontinuous carbon fibre reinforced thermoplastic composites with high fibre efficiency. *Addit. Manuf.* **2025**, *97*, 104587. [[CrossRef](#)]
37. Pacella, M.; Papa, A.; Papadia, G.; Fedeli, E. A Scalable Framework for Sensor Data Ingestion and Real-Time Processing in Cloud Manufacturing. *Algorithms* **2025**, *18*, 22. [[CrossRef](#)]
38. AlSalem, T.S.; Almaiah, M.A.; Lutfi, A. Cybersecurity Risk Analysis in the IoT: A Systematic Review. *Electronics* **2023**, *12*, 3958. [[CrossRef](#)]
39. Bakas, G.; Bei, K.; Skaltsas, I.; Gkartzou, E.; Tsiokou, V.; Papatheodorou, A.; Karatza, A.; Koumoulos, E.P. Object detection: Custom trained models for quality monitoring of fused filament fabrication process. *Processes* **2022**, *10*, 2147. [[CrossRef](#)]
40. Badarinath, R.; Prabhu, V. Real-time sensing of output polymer flow temperature and volumetric flowrate in fused filament fabrication process. *Materials* **2022**, *15*, 618. [[CrossRef](#)]
41. Zhang, L.; Chen, X.; Zhou, W.; Cheng, T.; Chen, L.; Guo, Z.; Han, B.; Lu, L. Digital twins for additive manufacturing: A state-of-the-art review. *Appl. Sci.* **2020**, *10*, 8350. [[CrossRef](#)]
42. Frunzaverde, D.; Cojocar, V.; Ciubotariu, C.R.; Miclosina, C.O.; Ardeljan, D.D.; Ignat, E.F.; Marginean, G. The influence of the printing temperature and the filament color on the dimensional accuracy, tensile strength, and friction performance of FFF-printed PLA specimens. *Polymers* **2022**, *14*, 1978. [[CrossRef](#)] [[PubMed](#)]
43. Bergaliyeva, S.; Sales, D.L.; Delgado, F.J.; Bolegenova, S.; Molina, S.I. Effect of thermal and hydrothermal accelerated aging on 3D printed polylactic acid. *Polymers* **2022**, *14*, 5256. [[CrossRef](#)] [[PubMed](#)]
44. Orellana-Barrasa, J.; Tarancón, S.; Pastor, J.Y. Effects of Accelerating the Ageing of 1D PLA Filaments after Fused Filament Fabrication. *Polymers* **2022**, *15*, 69. [[CrossRef](#)] [[PubMed](#)]
45. Calderon, G.; del Campo, G.; Saavedra, E.; Santamaría, A. Monitoring framework for the performance evaluation of an IoT platform with Elasticsearch and Apache Kafka. *Inf. Syst. Front.* **2023**, *26*, 2373–2389. [[CrossRef](#)]
46. Al-Ali, A.; Gupta, R.; Zualkernan, I.; Das, S.K. Role of IoT technologies in big data management systems: A review and Smart Grid case study. *Pervasive Mob. Comput.* **2024**, *101*, 101905. [[CrossRef](#)]
47. Rajpurohit, A.M.; Kumar, P.; Kumar, R.R.; Kumar, R. A Review on Apache Spark. *Kilby* **2023**, *100*, 7. [[CrossRef](#)]
48. Mirza, N.M.; Ali, A.; Ishak, M.K. The scheduling techniques in the Hadoop and Spark of smart cities environment: A systematic review. *Bull. Electr. Eng. Inform.* **2024**, *13*, 453–464. [[CrossRef](#)]
49. Tang, R.; Aridas, N.K.; Talip, M.S.A. Design of a data processing method for the farmland environmental monitoring based on improved Spark components. *Front. Big Data* **2023**, *6*, 1282352. [[CrossRef](#)]
50. Atitallah, S.B.; Driss, M.; Ghzela, H.B. Microservices for data analytics in IoT applications: Current solutions, open challenges, and future research directions. *Procedia Comput. Sci.* **2022**, *207*, 3938–3947. [[CrossRef](#)]
51. El Akhdar, A.; Baidada, C.; Kartit, A.; Hanine, M.; García, C.O.; Lara, R.G.; Ashraf, I. Exploring the Potential of Microservices in Internet of Things: A Systematic Review of Security and Prospects. *Sensors* **2024**, *24*, 6771. [[CrossRef](#)]
52. Bellini, A.; Güçeri, S.; Bertoldi, M. Liquefier dynamics in fused deposition. *J. Manuf. Sci. Eng.* **2004**, *126*, 237–246. [[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.