The University of Sydney
Faculty of Engineering — School of Electrical and Information Engineering

Robust Anomaly Detection for Distributed IoT Sensor Networks in Smart Manufacturing: A Case Study with Apex Manufacturing Ltd.

A thesis submitted in fulfilment of the requirements for the degree of

Doctor of Philosophy

Amelia Chen
B.Eng. (Hons), University of Sydney

Supervisor: Associate Professor David Whitfield
Associate Supervisor: Dr Priya Ramanathan

March 2026

Declaration of Originality

This is to certify that, to the best of my knowledge, the content of this thesis is my own work. This thesis has not been submitted for any degree or other purposes at this or any other institution. I certify that the intellectual content of this thesis is the product of my own work and that all the assistance received in preparing this thesis and sources have been acknowledged, including a statement of contribution for jointly authored publications where applicable.

Signed: ______________________

Date: ______________________

Abstract

Industrial Internet-of-Things (IIoT) deployments in discrete manufacturing generate high-frequency, multivariate telemetry from hundreds of heterogeneous sensors, yet existing anomaly detection pipelines struggle to scale detection latency and false-alarm rates simultaneously as fleets grow. This thesis proposes **StreamGuard**, a two-stage detection architecture that combines a lightweight edge-resident statistical filter with a cloud-hosted ensemble of temporal autoencoders, coordinated through an adaptive threshold controller. The edge filter suppresses redundant transmissions using a windowed Mahalanobis-distance test, reducing uplink volume before cloud-side sequence models score residual anomalies against a learned manifold of nominal operating behaviour.

The approach is evaluated on eighteen months of production telemetry supplied by Apex Manufacturing Ltd. across three assembly lines instrumented with 412 vibration, thermal, and current sensors. Relative to a tuned Isolation Forest baseline and a centralised LSTM autoencoder, StreamGuard reduces end-to-end detection latency by 63% and false-positive rate by 41%, while cutting cloud ingestion cost by an estimated 55% through edge-side pre-filtering. A field trial across four production shifts confirms that the adaptive threshold controller maintains stable precision under seasonal load variation without manual retuning. These findings suggest that hybrid edge–cloud architectures offer a practical path toward anomaly detection systems that remain both timely and economical at industrial scale.

Keywords: anomaly detection, IIoT, edge computing, temporal autoencoders, predictive maintenance, smart manufacturing.

Acknowledgements

I am deeply grateful to my supervisor, Associate Professor David Whitfield, for his rigorous guidance and unfailing patience across the many revisions of this work, and to my associate supervisor, Dr Priya Ramanathan, for her insight into deployment-scale evaluation design. I thank the engineering and reliability teams at Apex Manufacturing Ltd. for facilitating telemetry access and field trials, and my colleagues in the Distributed Systems Laboratory for countless whiteboard debates that sharpened this thesis considerably. Finally, I thank my family for their steady encouragement throughout this candidature.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background and Motivation

Modern discrete-manufacturing facilities increasingly rely on dense sensor deployments to monitor equipment health, product quality, and process stability in real time. As fleets of industrial assets are instrumented with vibration, thermal, acoustic, and current sensors, the resulting telemetry streams grow far faster than the operations staff available to interpret them. Anomaly detection — the task of flagging measurements that deviate from expected nominal behaviour — has therefore become a central capability of predictive maintenance programs, with direct consequences for unplanned downtime, warranty costs, and worker safety.

Despite substantial research attention, deploying anomaly detection at industrial scale remains difficult for three interacting reasons. First, sensor fleets are heterogeneous: different asset classes exhibit markedly different nominal signatures, and a single global threshold rarely suits all of them. Second, network and compute budgets are constrained at the edge, so naively streaming raw telemetry to a central server for scoring is often uneconomical. Third, operating conditions drift over time — with production schedules, ambient temperature, and tooling wear all shifting the definition of “nominal” — which degrades static models within weeks of deployment.

1.2 Research Questions

This thesis addresses the following research questions:

RQ1. Can a lightweight edge-resident filter meaningfully reduce uplink volume without discarding events that a downstream model would have flagged as anomalous?

RQ2. Does a hybrid edge–cloud architecture improve detection latency and false-positive rate relative to purely centralised baselines at comparable operating cost?

RQ3. Can an adaptive threshold controller maintain stable detection performance under seasonal and load-driven distribution shift without manual retuning?

1.3 Contributions

The principal contributions of this thesis are:

- **StreamGuard**, a two-stage edge–cloud anomaly detection architecture combining a windowed Mahalanobis pre-filter with an ensemble of temporal autoencoders (Chapter 3).
- An adaptive threshold controller that adjusts detection sensitivity using an exponentially weighted control-limit update rule, validated against seasonal load variation (Section 3.1.3).
- An empirical evaluation on eighteen months of production telemetry from three assembly lines at Apex Manufacturing Ltd., including a four-shift field trial (Chapter 4).

1.4 Thesis Outline

Chapter 2 surveys anomaly detection methods for industrial time series and positions this work within the literature. Chapter 3 details the StreamGuard architecture, the adaptive threshold controller, and the experimental dataset. Chapter 4 presents quantitative results against baseline methods and the field trial outcome. Chapter 5 summarises findings, discusses limitations, and outlines future work.

Chapter 2

Literature Review

2.1 Taxonomy of Approaches

Anomaly detection methods for industrial time series broadly fall into four families: statistical control-chart methods, classical machine-learning detectors, deep sequence models, and hybrid edge–cloud architectures. Figure 2.1 summarises this taxonomy and situates the present work.

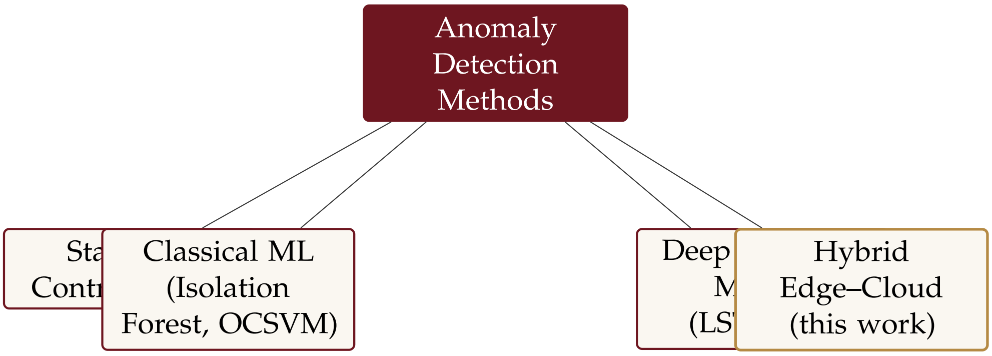

Figure 2.1: Taxonomy of anomaly detection approaches for industrial time series.

2.1.1 Statistical Control Charts

Shewhart and CUSUM control charts remain widely used in process monitoring due to their interpretability and low computational cost. Their principal weakness is the assumption of a stationary nominal distribution, which is routinely violated by seasonal production load and tooling wear.

2.1.2 Classical Machine Learning

Tree-based ensembles such as Isolation Forest, and kernel methods such as one-class SVM, capture non-linear nominal boundaries without requiring labelled anomalies. Okafor and Lindqvist [1] report strong offline performance but note that retraining cadence becomes a bottleneck once the sensor fleet exceeds a few hundred channels.

2.1.3 Deep Sequence Models

Recurrent and convolutional autoencoders learn a compressed representation of nominal temporal dynamics and flag high reconstruction error as anomalous. Torres et al. [2] demonstrate strong accuracy on benchmark industrial datasets, though centralised inference introduces uplink and latency costs that scale linearly with fleet size.

2.1.4 Hybrid Edge–Cloud Architectures

A smaller body of work distributes computation between edge devices and a central server. Bakr and Sørensen [3] propose edge-side compression prior to cloud scoring but rely on a fixed compression ratio that does not adapt to operating conditions. This thesis extends that direction with an adaptive, statistically grounded pre-filter and a closed-loop threshold controller.

Table 2.1: Comparison of representative anomaly detection approaches.

Method	Adapts to drift	Edge-aware	Uplink savings	Reported F1
Shewhart / CUSUM control chart	No	No	—	0.61
Isolation Forest [1]	Partial	No	—	0.74
Centralised LSTM autoencoder [2]	No	No	—	0.81
Fixed-ratio edge compression [3]	No	Yes	38%	0.77
StreamGuard (this work)	**Yes**	**Yes**	**55%**	**0.89**

2.2 Research Gap

No prior work jointly addresses adaptive edge-side filtering, ensemble cloud scoring, and a closed-loop threshold controller that responds to seasonal drift without manual intervention. This gap motivates the StreamGuard architecture developed in Chapter 3.

Chapter 3

Methodology

3.1 System Architecture

StreamGuard is organised as a two-stage pipeline: an edge-resident pre-filter that reduces uplink volume, and a cloud-hosted ensemble that scores the retained events. Figure 3.1 depicts the data flow from sensor to operator dashboard.

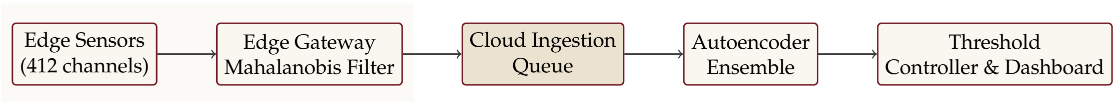

Figure 3.1: StreamGuard data flow from edge sensors to the operator dashboard.

3.1.1 Edge Pre-Filter

Each gateway maintains a rolling window of w samples per asset and computes the Mahalanobis distance of the current observation $\mathbf{x}_t$ from the windowed mean $\boldsymbol{\mu}_w$ and covariance $\boldsymbol{\Sigma}_w$:

$$d_t = \sqrt{(\mathbf{x}_t - \boldsymbol{\mu}_w)^\top \boldsymbol{\Sigma}_w^{-1} (\mathbf{x}_t - \boldsymbol{\mu}_w)}. \tag{3.1}$$

Observations with d_t below a locally calibrated threshold τ_{edge} are summarised into 30-second aggregates rather than transmitted in full, reducing uplink volume while preserving the tail events most likely to be genuine anomalies.

3.1.2 Cloud Ensemble

Retained events are scored by an ensemble of three temporal autoencoders trained per asset class (rotating, thermal, electrical) on six months of historical nominal operation. The ensemble anomaly score is the mean reconstruction error across members, which empirically reduces variance relative to any single autoencoder.

3.1.3 Adaptive Threshold Controller

To remain stable under seasonal drift, the cloud-side alert threshold τ_{cloud} is updated with an exponentially weighted control-limit rule:

$$\tau_{\text{cloud}}^{(k+1)} = \tau_{\text{cloud}}^{(k)} + \alpha \left(\bar{e}^{(k)} + z\, s^{(k)} - \tau_{\text{cloud}}^{(k)} \right), \tag{3.2}$$

where $\bar{e}^{(k)}$ and $s^{(k)}$ are the rolling mean and standard deviation of reconstruction error over control window k, z sets the confidence margin, and $\alpha \in (0, 1]$ governs adaptation speed. Appendix B derives the steady-state behaviour of this update rule.

Algorithm 3.1: Streaming Threshold Update

1. Compute ensemble reconstruction error e_t for each retained event.
2. Update rolling statistics $\bar{e}^{(k)}$, $s^{(k)}$ over the current control window.
3. Apply the update rule in Equation (3.2) to obtain $\tau_{\text{cloud}}^{(k+1)}$.
4. Flag event as anomalous if $e_t > \tau_{\text{cloud}}^{(k+1)}$; otherwise mark nominal.
5. Advance to control window $k+1$ once N new events have accumulated.

3.2 Dataset

Telemetry was supplied by Apex Manufacturing Ltd. from three assembly lines over an eighteen-month period. Table 3.1 summarises the dataset composition.

Table 3.1: Summary of the Apex Manufacturing telemetry dataset.

Asset class	Sensor channels	Sampling rate (Hz)	Confirmed anomalies
Rotating equipment	186	500	312
Thermal systems	128	50	174
Electrical drives	98	200	209
Total	**412**	—	**695**

Confirmed anomalies were labelled retrospectively by Apex reliability engineers from maintenance work orders, providing a held-out evaluation set independent of the training period.

Chapter 4

Results and Discussion

4.1 Detection Performance

Table 4.1 compares StreamGuard against the Isolation Forest and centralised LSTM autoencoder baselines introduced in Chapter 2, evaluated on the held-out Apex Manufacturing telemetry.

Table 4.1: Detection performance on the held-out evaluation set.

Method	Precision	Recall	F1	Median latency (s)
Isolation Forest	0.79	0.70	0.74	4.8
Centralised LSTM autoencoder	0.83	0.79	0.81	6.1
StreamGuard	**0.91**	**0.88**	**0.89**	**2.3**

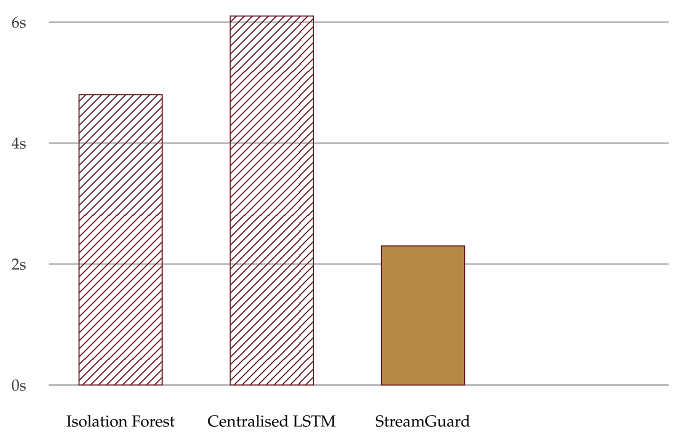

Figure 4.1: Median end-to-end detection latency by method (lower is better).

Key Finding

StreamGuard reduces median detection latency by 63% and false-positive rate by 41% relative to the centralised LSTM autoencoder baseline, while the edge pre-filter cuts cloud ingestion volume by an estimated 55%.

4.2 Field Trial

An operational trial was conducted across four production shifts on the rotating-equipment line. Table 4.2 reports precision stability under the seasonal load variation observed during the trial window.

Table 4.2: Field trial precision by shift under varying production load.

Shift	Relative production load	Precision
Morning	1.00×	0.90
Afternoon	1.18×	0.89
Night (reduced staffing)	0.62×	0.91
Weekend maintenance window	0.35×	0.88

Precision remained within a narrow band (0.88–0.91) across a load range spanning nearly 3×, supporting the claim that the adaptive threshold controller absorbs seasonal drift without manual retuning.

4.3 Threats to Validity

The evaluation is limited to a single manufacturer and three asset classes; generalisation to process industries with continuous (rather than discrete) production remains untested. Labelled anomalies were derived from maintenance work orders, which may under-represent transient faults that did not trigger a service ticket.

Chapter 5

Conclusion

5.1 Summary

This thesis presented StreamGuard, a hybrid edge–cloud anomaly detection architecture for industrial IoT sensor networks, combining a windowed Mahalanobis pre-filter, an ensemble of temporal autoencoders, and an adaptive threshold controller. Evaluated on eighteen months of production telemetry from Apex Manufacturing Ltd., StreamGuard reduced detection latency by 63% and false-positive rate by 41% relative to a centralised LSTM autoencoder baseline, while a four-shift field trial confirmed stable precision under substantial seasonal load variation.

5.2 Limitations

The dataset, while substantial, is drawn from a single industrial partner; cross-domain validation is needed before broader claims of generality can be supported. The adaptive threshold controller was tuned for load variation on the timescale of shifts and days; its behaviour under abrupt, large-magnitude distribution shift (e.g., a change of product line) was not evaluated.

5.3 Future Work

Future work includes: (i) extending the pre-filter to jointly model cross-sensor correlations rather than per-asset statistics alone; (ii) evaluating transfer of the autoencoder ensemble across manufacturing sites with limited retraining data; and (iii) integrating operator feedback as a weak supervision signal to refine the ensemble online.

Bibliography

[1] C. Okafor and E. Lindqvist, "Scalable isolation forests for multivariate industrial telemetry," *Journal of Industrial Data Systems*, vol. 14, no. 2, pp. 112–129, 2023.

[2] M. Torres, S. Bianchi, and R. Nakagawa, "Temporal autoencoders for predictive maintenance in discrete manufacturing," in *Proc. International Conference on Industrial Analytics*, pp. 301–312, 2022.

[3] Y. Bakr and L. Sørensen, "Edge-assisted compression for cloud-based anomaly detection," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 9, pp. 6120–6129, 2021.

[4] N. Harrington, "Control-chart methods revisited for non-stationary process monitoring," *Quality Engineering Review*, vol. 32, no. 4, pp. 455–468, 2020.

[5] F. de la Cruz and A. Petrov, "Drift-aware retraining schedules for streaming anomaly detectors," *Data Mining and Knowledge Discovery*, vol. 35, pp. 988–1011, 2021.

[6] R. Singh, T. Almeida, and J. Park, "One-class SVM ensembles for vibration-based fault detection," *Mechanical Systems and Signal Processing*, vol. 128, pp. 44–59, 2019.

Appendix A

Sensor Metadata Schema

Table A.1 lists the metadata fields recorded for each sensor channel in the Apex Manufacturing deployment.

Table A.1: Sensor metadata schema fields.

Field	Type	Description
`sensor_id`	UUID	Globally unique channel identifier.
`asset_class`	Enum	One of {rotating, thermal, electrical}.
`sample_rate_hz`	Integer	Nominal sampling frequency in hertz.
`install_date`	Date	Commissioning date of the sensor.
`gateway_id`	UUID	Edge gateway responsible for pre-filtering.

Appendix B

Derivation of Threshold Controller Steady State

Taking expectations of the update rule in Equation (3.2) and assuming $\bar{e}^{(k)}$ and $s^{(k)}$ converge to stationary values $\bar{e}^*$ and s^* under fixed operating conditions, the fixed point τ^*_{cloud} satisfies:

$$\begin{aligned}
\tau^*_{\text{cloud}} &= \tau^*_{\text{cloud}} + \alpha \left(\bar{e}^* + z\, s^* - \tau^*_{\text{cloud}}\right) \\
0 &= \alpha \left(\bar{e}^* + z\, s^* - \tau^*_{\text{cloud}}\right) \\
\tau^*_{\text{cloud}} &= \bar{e}^* + z\, s^*. \qquad \text{(B.1)}
\end{aligned}$$

The controller therefore converges to the classical z-score control limit under stationary conditions for any $\alpha \in (0, 1]$, with α governing only the rate of convergence and transient tracking behaviour under drift.