Structured Degradation in Visible Light Positioning

Modeling and Compensation of Long-Term Degradation in RSS-Based VLP Systems

Jasper van Arkel
Supervisor: Dr. Qing Wang, Shun Zhuge, Bo Yang

EEMCS, Delft University of Technology, The Netherlands

A thesis submitted to the Faculty of Electrical Engineering, Mathematics and Computer Science
Delft University of Technology
in partial fulfilment of the requirements for the degree of
Bachelor of Computer Science and Engineering
June 21, 2026

Name of the student: Jasper van Arkel
Final project course: CSE3000 Research Project
Thesis committee: Dr. Qing Wang, Shun Zhuge, Bo Yang Mark Neerincx

***Abstract*—Visible Light Positioning (VLP) uses LEDs for accurate indoor localization. However, structured illumination drift caused by LED aging, optical contamination, thermal effects, blockages, and device failures can reduce the long-term accuracy of RSS-based VLP systems. This thesis investigates how this drift can be modeled and compensated for using lightweight algorithms suitable for microcontrollers.**

The proposed method combines scaling-based compensation for gradual degradation with anomaly detection for sudden degradation events such as broken LEDs. This method is tested through a long-term deployment simulation using the DenseVLC dataset and is also implemented on a Raspberry Pi Pico to assess embedded workability. The results show that VLP Systems suffer more errors over time, while degradation-aware compensation improves long-term robustness. However, embedded deployment introduces accuracy trade-offs due to quantization and memory constraints.

This data shows that modeling and compensating degradation mechanisms is important for reliable long-term VLP deployment, and that compensation methods needs to account for both gradual and sudden changes in received signal strength.

I. Introduction

Visible Light Positioning (VLP) is a technique that uses LEDs for accurate indoor localization[5]. Traditional localization methods, such as GPS, struggle with accurate (indoor) localization, only being accurate up to several meters[6]. Indoor localization techniques such as WiFi, Radio-Frequency identification, and Ultra-Wide-band are often expensive due to the need to install Wi-Fi access points or transmitting equipment. These systems only reach up to 1-meter accuracy, whereas VLP systems have demonstrated centimeter-level accuracy[6]. Because of its higher accuracy and cheaper implementation costs, VLP represents a strong alternative. Multiple techniques have been used to implement VLP systems. The main variations of techniques include Angle-of-Arrival (AOA), Time of Arrival, Time-difference-of-arrival, and Received signal strength [6].

RSS-based systems require fingerprint measurements. This, however, is prone to LED degradation and other types of distortion, including structured illumination drift, in which LED intensity changes occur in a spatially correlated manner due to LED aging, failure, or surrounding conditions. creating a need for regular recalibration. Recent work proposed an online solution that automatically adapts to an aging framework and degrading LEDs without requiring manual recalibration, making the system suitable and robust for long-term deployment.

This thesis addresses the following research question:

How can structured illumination drift in RSS-based Visible Light Positioning systems be modeled and compensated for using lightweight algorithms suitable for microcontrollers?

This research question is divided into four subquestions:

1) Which degradation mechanisms contribute to structured illumination drift, and how can they be modeled?
2) How can gradual and sudden illumination drift be compensated for?
3) How can the compensation method be implemented on a resource-constrained microcontroller?
4) How robust is the method under combined degradation during long-term deployment?

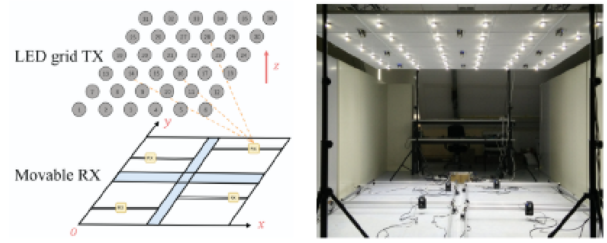

(a) Schematic (b) Testbed

Fig. 1. The DenseVLC testbed [1]

II. Background

This section introduces the background knowledge required for this research. First, RSS-based Visible Light Positioning is explained. Next, the DenseVLC dataset used in the experiments is described. Then, LED degradation and finally an existing per-LED scaling method is introduced to compensate for age degradation over time.

A. RSS-based Visual light positioning

Received Signal Strength (RSS)-based Visible Light Positioning systems can be implemented using cost-efficient hardware. In RSS systems, multiple LEDs act as transmitters, and a camera or photodiode measures the light intensity received from these LEDs. Photodiodes are preferred in many applications because they are less expensive and require fewer computational resources than cameras.

RSS-based positioning is implemented using fingerprinting. First measurements of light intensities are being made from known locations and stored in a fingerprint database. During deployment, measured RSS values are fed into a positioning model trained on fingerprints to estimate the receiver's location. This causes the density and quality of the collected fingerprints to directly influence the positioning performance [5].

B. Dataset

In VLP a proper fingerprint database is of great importance for the accuracy of the system. For this research, we use an indoor position dataset, which is collected on the DenseVLC Testbed [1]. This setup consists of 36 ceiling-mounted LEDs and four movable sensors, as shown in Fig 1. Each sample consists of 36 RSS measurements and a ground truth position. The distance between different measurements is one cm in both the x and y direction.

C. Degradation

Degradation in VLP Systems is a known issue. The effect from age degradation alone has been shown to cause up to $19.45cm$ errors using Multi-layer perceptrons in a simulation environment to simulate a 100000 hour deployment. The aging of LED is defined by the TM-21 standard[10], which is an industry agreement in which they state the LED age degradation is measured as:

$$\Phi(t) = \beta e^{-\alpha t} \tag{1}$$

Where Φ is the lumen maintenance β is the pre-factor, and α is the decay parameter. This α is different for each type of LED, differs even between batches of LEDs made, and is temperature dependent.

D. Per LED Scaling

Previous research proposed an online per LED scaling algorithm that dynamically learns the gradual degradation of the LED's assuring accurate positioning regardless of environmental and infrastructure changes. This method created a huge improvement in the long-term reliability of the VLP Systems, as this decreased the average error while using the multi-layer perceptron model to only $1.49cm$ and even to $0.49cm$ using WOKNN + online calibration after a 100000 hour deployment

III. Literature Review and Related Work

This chapter reviews the existing literature on the long-term reliability of RSS-based visual light positioning systems. Particular attention is given to mechanisms that degrade the transmitted or received signal strength, thereby reducing the VLP's performance. The chapter concludes by identifying the research gap this thesis addresses.

A. Types of LED Degradation

The existing Literature does not present a single unified classification of degradation types in a VLP System. Instead, different studies focus on specific physical, optical, or thermal processes. This section examines degradation mechanisms mentioned in multiple studies and describes their importance to VLP systems. The identified types of degradation are:

- LED aging [2]
- LED dirt [8]
- Sensor dirt [9]
- Blockages [11]
- Failures [3]
- Permanent damage [3]
- Thermal droop [7]

Not all of these mechanisms have been studied in the context of VLP. However, they are included as they either affect the emitted light or the received signal in VLP Systems.

B. Modeling LED Degradation

For each degradation mechanism, its causes and relationship with the measured light in VLP have been analyzed.

1) LED Aging: LED aging has a prediction standard published by the IESNA, TM-21[2]. This is used to determine the operating lifetime of LED light sources based on data collected during the IES LM-80-08 test. For Lumen maintenance, Equation 1 is used. α is temperature-dependent and can be estimated using the following equation:

$$\alpha = AI^n \exp(-\frac{E_a}{kT_j}) \tag{2}$$

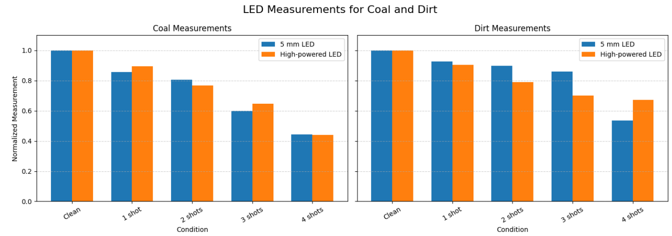

Fig. 2. Re-plotted measurements based on data reported in [8]. The original charts were re-visualized for readability. The chart shows Dirt and Coal in two separate plots for both a $5mm$ LED diode and a high-power LED lamp.

Where α is the degradation factor, A is the exponential factor, and n is the current acceleration exponent, which are both unique to the LED type itself. E_a represents the activation energy, k is the Boltzmann constant and T_j is the absolute juncture temperature in kelvin.

2) Optical Surface Contamination: Contamination buildup can decrease the amount of light transmitted and measured throught the measurement system. Because this contamination happens on both LED and the Sensor both types of contamination are both explained as optical surface contamination together.

The lumen output for LED lightning is affected by contamination on the surface of the LED. The decrease in light output is dependent on the type of contamination as well as the amount of dirt on the lens and the shape of the lens [8]. Water and fingerprints leave no measurable difference in lumen output. But as shown in Fig 2, increasing the amount of soil and coal on the LED's lens decreases the light output greatly.

While this study on LED contamination show that contamination decreases the lumen output. A separate deployment study provides direct measurements of light throughput through a contaminated medium over time [9]. It shows that the contamination is depended on the angle of the medium as well as real in world happenings such as spiderwebs being made and Sahara dust storms. But the measurements shown in Fig 3 show that after a 730-day deployment, measurements for a medium on 90 degrees (sensor) the measured light degraded by $1 - 2,5\%$, and for the surface on 180 degrees (ceiling-mounted lamps) it showed a decrease of $8 - 15\%$, but this was mainly due to a spiderweb retaining more dirtiness

3) Blockages, Failures, and Damage: In addition to gradual degradation, some mechanisms occur suddenly, such as blockages, failures, and physical damage. These mechanisms can cause a faster and more abrupt reduction in light output than the previously discussed degradation mechanisms. Blockages can obstruct the line of sight between the LED and the receiver, thereby reducing the received signal strength. They have been shown to decrease the positioning accuracy of an RSS-based VLP system by up to 90% [11].

4) Thermal Droop: A source of fluctuating degradation is thermal droop. Thermal droop is the degradation of the light output of LEDs due to a rising junction temperature. This can be caused by multiple factors, including, but not limited to, increased non-radiative recombination, carrier leakage, and temperature-dependent changes in carrier transport and recombination processes within the semiconductor material [7]. A temperature increase can reduce LED luminous output by approximately 0.2% to 1% per degree Celsius [4]. Therefore, fluctuations in junction temperature can result in corresponding fluctuations in the emitted light output.

C. Compensation Methods for LED Degradation

In recent work [5], the effect of aging on the VLP system was looked at, and an online learning algorithm was proposed, which dynamically updated the LED's RSS values. Instead of requiring manual adjusting of the full fingerprint database, the method adapts the measured RSS values to compensate for gradual changes in LED output.

This approach was shown to significantly improve long-term positioning accuracy. With simulated LED aging, the positioning error was reduced from 19.45 cm without compensation to 0.49 cm when using WOKNN with online calibration. This demonstrates that degradation-aware compensation can substantially improve the robustness of RSS-based VLP systems during long-term deployment.

However, this method mainly focuses on LED aging. It remains unclear whether similar compensation strategies are sufficient when other degradation mechanisms, such as contamination, thermal droop, blockages, or LED failures, occur either individually or in combination.

D. Effects on VLP Systems

The degradation mechanisms discussed in this section have important implications for RSS-based VLP. The RSS-based localization assumes that the fingerprints remain consistent over time, the changes caused by the mentioned degradations can severely reduce accuracy.

Gradual degradation mechanisms such as LED aging and optical contamination cause a growing difference between the RSS fingerprints used during training and the measurements observed during deployment. As this drift increases, the positioning model becomes less of a representation of the real environment, causing larger localization errors.

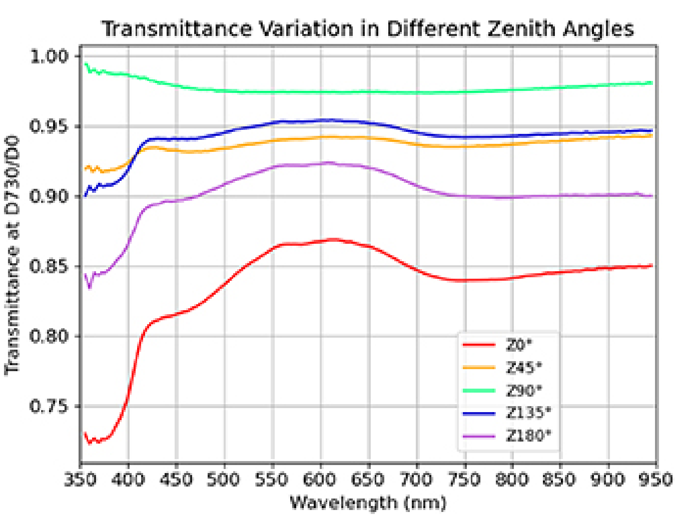

Fig. 3. Transmittance variation measured after a 730 day deployment.

Furthermore, thermal droop introduces temporary changes in signal strength, while blockages, failures, and permanent damage can cause sudden and significant drops in the received signal. These abrupt changes may result in large positioning errors if they are not detected and handled appropriately.

Together, these degradation mechanisms highlight a key challenge for long-term VLP deployment: staying accurate regardless of changing environmental and hardware conditions. Compensation methods that can adapt to both gradual and sudden degradation effects are required to ensure reliable and robust positioning performance without regular manual recalibration.

E. Research Gap

While the effect of single degradation mechanisms has been researched [5] [11], the summation of all of these mechanisms has not been made yet. Furthermore, there has not been research on how the combination of these factors affects the accuracy of RSS-based VLP. In real deployments, multiple degradation mechanisms may occur simultaneously. For example, LED aging can happen at the same time as optical contamination, thermal effects, partial blockages, or hardware failures. The combination of these mechanisms may affect received signal strength measurements and affect the VLP system in ways that cannot be accurately predicted by studying each mechanism independently. Furthermore, existing compensation methods are generally designed to address a specific degradation source, most notably LED aging. It remains unclear how effective such methods are when multiple degradation mechanisms occur concurrently. Therefore, there is a need for an evaluation of the combined effects of degradation mechanisms and the development of compensation strategies that improve the robustness of RSS-based VLP systems during long-term deployment. This thesis targets this gap by identifying and modeling multiple degradation mechanisms, evaluating their impact on positioning accuracy, and investigating compensation methods that can maintain reliable and robust localization performance over longer deployment periods.

IV. Methodology

This chapter describes the research methodology used to investigate the impact of long-term degradation on RSS-based VLP systems and to evaluate compensation methods that improve system robustness during long-term deployment.

A. Research Approach and Workflow

This research follows a simulation-based experimental method. First, degradation mechanisms that can affect VLP systems during long-term deployment are identified through a literature review. For these degradation mechanisms, mathematical models are created that can be applied to RSS measurements. Lastly, compensation methods are implemented and evaluated under various degradation scenarios.

The methodology addresses the four subquestions by modeling relevant degradation mechanisms, developing compensation methods for gradual and sudden drift, implementing the method on a microcontroller, and evaluating its robustness during simulated long-term deployment.

B. Degradation Modeling

Based on the degradation mechanisms identified in Section III, four degradation categories are incorporated into the simulation framework: LED aging, optical surface contamination, thermal droop, and sudden degradation events. This includes all identified mechanisms. The original DenseVLC RSS measurements are treated as measurements from an undegraded environment. During the simulation, degradation factors are applied to these measurements to represent changes in emitted or received light intensity over time. When multiple degradation mechanisms occur simultaneously, their corresponding scaling factors are combined before the degraded RSS vector is passed to the positioning model.

1) LED Aging: LED aging is modeled as a gradual and continuous decrease in emitted light intensity over time. The degradation follows the lumen maintenance model discussed in Section III, where each LED experiences a reduction in output according to its degradation parameters. This is assumed to be nearly equal as the assumption is made that the same type of LEDs from the same batch are being used. Small variations between individual LEDs are included to represent manufacturing differences and environmental effects. The resulting aging factor is updated throughout the simulated deployment and applied independently to the RSS contribution of each LED.

2) Optical Surface Contamination: Optical surface contamination represents the accumulation of dirt and dust on either LED lenses or sensor surfaces. This contamination is modeled as a scaling factor applied to the measured RSS values. In addition, cleaning events are simulated by restoring the contamination scaling factors to their undegraded values. Other degradation effects, such as LED aging or permanent damage, remain unchanged after cleaning.

3) Thermal Droop: Thermal droop is modeled as a temporary reduction in light output caused by an increase in LED junction temperature. Unlike LED aging, thermal droop is a temporary, reversible reduction in output and therefore introduces short-term fluctuations in the received signal strength rather than permanent degradation.

Since this thesis focuses on indoor VLP deployments, room temperature is assumed to be the normal operating condition. Thus, thermal droop is modeled as an occasional effect that occurs during periods of elevated ambient temperature or reduced cooling efficiency. These events result in a temporary increase in LED junction temperature, causing a corresponding reduction in emitted light intensity. Once normal operating conditions are restored, the LED output returns to its nominal level.

4) Blockages, Failures and Damages: Sudden degradation mechanisms are represented by blockages, failures and permanent damage. Blockages are modeled by temporarily reducing or removing the contribution of one or more LEDs. LED failures are represented by a complete loss of signal from the affected LED, while permanent damage is modeled as a sudden and lasting reduction in emitted light intensity.

C. Proposed Compensation Method

This thesis combines two types of compensation methods for the identified degradation mechanisms. For gradual degradation mechanisms, such as LED aging, optical contamination, and thermal droop, a scaling-based compensation approach is used. These mechanisms reduce the emitted as well as received light intensity in a way that can be approximated by a degradation factor. Therefore, the affected RSS values are corrected by applying scaling factors derived from the corresponding degradation models.

For sudden degradation mechanisms, including blockages, LED failures, and permanent damage, a different approach is required. These mechanisms can cause abrupt and significant changes in received signal strength. To detect such events, the compensated RSS measurements are compared with expected RSS values obtained from the fingerprint database.

First, the scaling-corrected RSS vector is used to produce an initial position estimate. A KNN algorithm then selects the (K) reference points closest to this estimated position and combines their associated RSS measurements to construct an expected RSS vector, making use of a weighted average. Measured RSS values that differ from their corresponding expected values by more than a predefined threshold are flagged as anomalous. These anomalous values are then replaced with their expected values before the final position estimate is calculated.

The complete method, therefore, consists of two consecutive stages as shown in Algorithm 1. First, scaling-based compensation corrects predictable reductions in RSS. Second, anomaly detection detects and corrects sudden signal changes that cannot be addressed using the degradation models alone. By combining these strategies, the proposed method aims to improve the robustness of RSS-based VLP systems under changing conditions during long-term deployment.

D. Evaluation Strategy

The effectiveness of the proposed compensation method is evaluated under a combined long-term degradation scenario in which all modeled degradation mechanisms occur during the same simulated deployment. These mechanisms include LED aging, optical surface contamination, thermal droop, blockages, LED failures, and permanent damage. This combined scenario is intended to represent a realistic deployment in which several environmental and hardware-related effects may influence the received signal strength simultaneously. The proposed method is compared with baseline positioning methods that do not use the complete degradation-aware compensation approach. At each simulated time step, the average positioning

Algorithm 1 Degradation-aware RSS positioning with anomaly correction

Require: Measured RSS vector $\mathbf{r} \in R^N$, luminaire age vector $\mathbf{A}$, temperature vector $\mathbf{T}$, dust accumulation vector $\mathbf{D}_u$, fingerprint database $\mathcal{D}$, threshold $\tau = 0.30$
Ensure: Estimated position $\hat{\mathbf{p}}$

Step 1: Compute gradual degradation factors
for $i = 1$ to N **do**
  $s_{a,i} \leftarrow f_a(A_i)$
  $s_{t,i} \leftarrow f_t(T_i)$
  $s_{d,i} \leftarrow f_d(D_{u,i})$
end for

Step 2: Compensate gradual degradation
for $i = 1$ to N **do**
  $s_i \leftarrow s_{a,i} \cdot s_{t,i} \cdot s_{d,i}$
  $r_i^c \leftarrow r_i \div s_i$
end for

Step 3: Initial RSS-based position estimation
$\hat{\mathbf{p}}_0 \leftarrow \mathrm{RSSPosition}(\mathbf{r}^c, \mathcal{D})$

Step 4: Estimate expected RSS values using KNN
$\mathbf{r}^e \leftarrow \mathrm{KNNExpectedRSS}(\hat{\mathbf{p}}_0, \mathcal{D})$

Step 5: Detect and correct anomalous LEDs
for $i = 1$ to N **do**
  if $r_i^c < (1 - \tau) \cdot r_i^e$ **then**
    $r_i^c \leftarrow r_i^e$
  end if
end for

Step 6: Final RSS-based position estimation
$\hat{\mathbf{p}} \leftarrow \mathrm{RSSPosition}(\mathbf{r}^c, \mathcal{D})$
return $\hat{\mathbf{p}}$

error is calculated over all evaluated samples. This produces an error trajectory that shows how positioning performance changes throughout the deployment period. The evaluation considers both positioning accuracy and stability over time. A robust method should achieve a low average positioning error while also preventing large temporary increases in error. Therefore, the error trajectories are examined for sudden peaks caused by stronger or abrupt degradation events. The proposed method is considered more robust when it maintains a relatively low error and limits large fluctuations throughout the deployment. Because all degradation mechanisms are evaluated together, the experiment measures the overall robustness of the methods under combined degradation rather than the individual effect of each mechanism. The results therefore show how well each method performs under the complete simulated deployment scenario, but they do not isolate the contribution of individual degradation sources.

V. Embedded Deployment and Simulation Setup

This section describes the experimental configuration used to evaluate the proposed degradation-aware positioning method. It presents the software and hardware environment, simulation parameters, KNN and anomaly-detection settings, evaluation metrics, and measures taken to support reproducibility.

A. Software Stack

The simulation framework built for this thesis is implemented in Python 3.12. Python was used in preparing the DenseVlC RSS measurements, applying the degradation, running the long-term deployment simulation, and evaluating the positioning errors.

The embedded part of this thesis uses a two-part software stack. The computationally expensive operations, such as training, degrading simulation, plotting, and evaluating, are performed on a host computer, and a Raspberry Pi Pico is used for real-time embedded inference. The software for the Pico is written is C++ using the Raspberry Pi Pico SDK and TensorFlow Lite Micro. The software for the host side of the embedded experiment is written in Python 3.12. During deployment, the host-side generates degraded sensor measurements. Each degraded sample is transmitted to the Pico, which performs embedded inference and returns a prediction. The host code then compares the Pico prediction with the actual position and records the positioning error. This allows the embedded model to be evaluated under long-term degradation conditions while keeping plotting and statistical analysis on the host machine.

B. Simulation Setup

For gradual degradation mechanisms such as aging, contamination, and thermal droop, scaling factors are applied to compensate for the expected reduction in RSS values. These scaling factors are derived from the corresponding degradation models and updated throughout the simulation.

For sudden degradation mechanisms such as blockages and failures, an anomaly detection approach is used. The anomaly detection algorithm monitors incoming RSS measurements and identifies signals that differ significantly from their expected behavior. Detected anomalies are being replaced with expected values before position estimation is performed. For the KNN anomaly detection, K = 1000. The anomaly threshold, $\tau = 0.30$, has been chosen because it is the value used in the TM-21 to define LED failure.

C. Evaluation Metrics

The positioning methods are evaluated using their error trajectories over the simulated deployment period. At each time step, the average Euclidean positioning error is calculated over all evaluated samples. This shows how positioning accuracy changes as degradation progresses and makes it possible to assess whether a method remains accurate throughout long-term deployment. The overall mean positioning error represents general accuracy, while the maximum positioning error

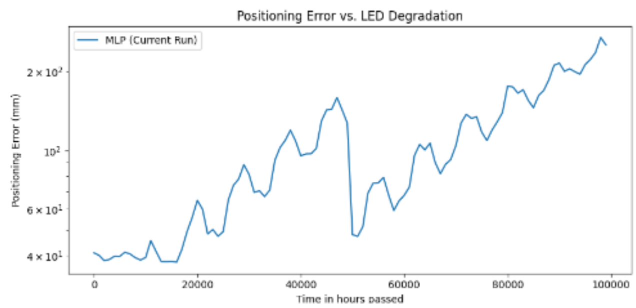

Fig. 4. Errors of MLP over 100000 hour deployment

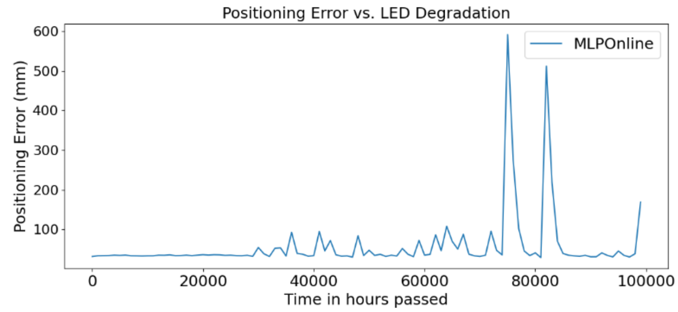

Fig. 5. Errors of MLP with online recalibration proposed by Li [5]

identifies severe temporary failures or error peaks. A robust method should therefore maintain both a low average error and a limited maximum error over the complete deployment period. Since embedded deployment is an important objective of this research, the inference runtime is also measured on the Raspberry Pi Pico. This metric is used to evaluate whether the proposed method can operate efficiently on resource-constrained hardware.

D. Reproducibility

To achieve reproducibility, all experiments use predefined random number generator seeds. For all experiments in this study, the used seed was ”42”. This guarantees that identical degradation events and simulation conditions can be reproduced across multiple experimental runs. The implementation can be found in the following GitHub repository: project GitHub repository.[1]

VI. Results

This section presents the results of the experiments conducted in this research. The goal of these experiments is to evaluate the effect of long-term degradation on RSS-based VLP systems and the effect of the proposed compensation algorithm. First, the effect of degradation is tested on existing methods. Afterwards, the proposed degradation-aware method is compared against these existing methods. Finally, the embedded implementation for the Raspberry Pi Pico is being evaluated.

A. Robustness During Long-Term Deployment

The long-term use analysis evaluates how positioning accuracy changes during a simulated deployment of 100000 hours. The experiments compare the positioning error of two baseline MLP models, an existing online calibration method, and the proposed degradation-aware compensation method. The goal of this analysis is to determine whether the proposed method can reduce the effect of long-term degradation on positioning accuracy.

1) Baseline and Online Calibration: The graph in Figure 4 shows the position ing error of the baseline MLP model during the simulated 100000- hour deployment. As shown in the graph, the baseline model's performance declines significantly over the deployment period. The positioning error increases over time, showing that the baseline model becomes less accurate as the RSS measurements drift further from the original fingerprint database. A decrease in error can be observed around the middle of the deployment period; this corresponds to a simulated cleaning event, in which all accumulated dust contamination on both the LEDs and the sensor is removed. After this event, the error starts increasing again as the contamination increases again and the degradation continues, this shows that cleaning the system can temporarily improve positioning accuracy.

Figure 5 shows the performance of the degradation-aware online calibration method proposed by Li [5]. Compared with the uncompensated MLP baseline, the online calibration method maintains a more stable positioning error over the deployment period. The results of this measurement indicate the effectiveness of a degradation-aware method,

However, the method still shows several large error spikes during the deployment. These spikes start occurring when the degradation effects become stronger. After these spikes, the positioning error returns close to its previous level, showing that the method can adapt over time but may still be affected by sudden or combined large degradation effects.

2) Compensation Under Combined Drift: Figure 6 shows the positioning error of the proposed degradation-aware compensation method when combined with both the MLP and Random Forest models. The MLP-based version remains relatively stable throughout the deployment. The Random Forest version shows larger fluctuations, but achieves lower positioning error during parts of the simulation where degradation effects are less severe.

This result shows that the choice of positioning model influences the behavior of the proposed compensation method. The MLP model provides more stable performance over time, while the Random Forest model can achieve lower errors in some parts of the deployment but is more sensitive to changes in degradation.

[1] https://github.com/JaspervArkel/VLP_Under_Degradation.git

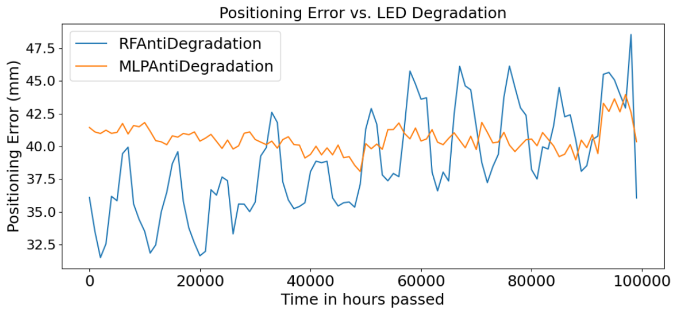

Fig. 6. Errors of MLP and Random Forest in combination with the proposed algorithm over 100000 hours of deployment

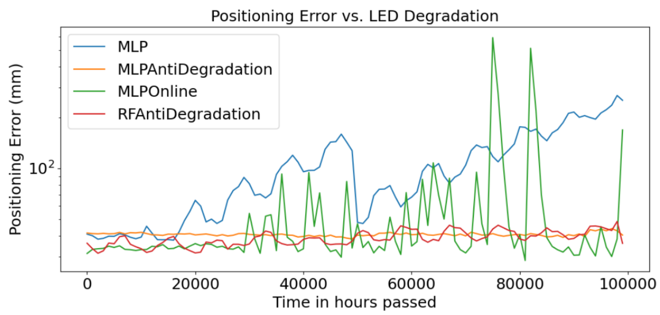

Fig. 7. Comparison between models without the anti-degradation method and Models combined with our proposed method

3) Microcontroller Implementation: The proposed method was also implemented for deployment on a Raspberry Pi Pico. For this implementation, the RSS measurement values were quantized to int8 values, and the measurement pool used for the KNN-based expected RSS estimation was reduced to fit the memory and computational constraints of the embedded device.

Figure 8 shows the positioning error over this implementation over the simulated period. This version starts with an error of approximately 40 mm, but the error increases more strongly over time in comparison with the desktop implementation. This shows that the method can be deployed on constrained hardware, but with reduced positioning accuracy compared the desktop version.

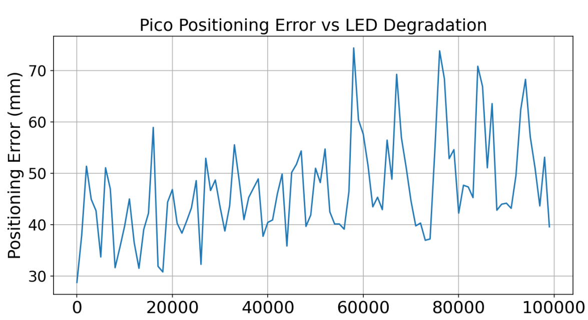

Fig. 8. Errors of MLP in combination with the proposed algorithm over 100000 hours of deployment on the RP Pico

TABLE I
CALCULATION-TIME MEASUREMENTS OF THE EMBEDDED IMPLEMENTATION ON THE RASPBERRY PI PICO.

Metric	**Calculation time**(μs)
Mean	3606
Median	3606
Minimum	3448
Maximum	4008
95th percentile	3657

The calculation time of the embedded implementation was measured on the Raspberry Pi Pico. As shown in Table I, the mean and median calculation times were both $3606\mu s$. The minimum observed calculation time was $3448\,\mu s$, while the maximum was $4008\,\mu s$. The 95th-percentile calculation time was $3657\,\mu s$, meaning that 95% of the calculations were completed within $3657\,\mu s$. These results show that the embedded implementation has relatively consistent execution times, with all measured calculations completing within at most $4008\,\mu s$.

VII. Responsible Research

This section discusses the responsible research aspects of this thesis, the main points discussed are ethical considerations, the reproducibility an the usage of generative AI.

A. Ethical Considerations

This research focuses on improving the long-term reliability of RSS-based Visible Light Positioning systems. Since VLP systems can be used for indoor tracking, one potential ethical concern is the use of a trained VLP Model to track people in indoor environments. This thesis uses an existing dataset and does not collect any personal data. Further training and deployment should take privacy into account. Users should be informed when such systems are active, and collected positioning data should be handled carefully to prevent misuse.

Another ethical consideration is the long-term reliability of positioning systems in safety-critical environments. If VLP is used in navigation, robotics, or tracking, degradation-related positioning errors could lead to incorrect decisions. Therefore, degradation-induced positioning errors must be detected and compensated. This thesis aims to improve positioning robustness and accuracy under changing conditions.

B. Reproducibility

The experiments in this thesis are simulation-based and use predefined random seeds. This makes it possible to reproduce the same degradation events, cleaning events, and positioning experiments over multiple runs. The degradation models, compensation method, and evaluation procedure are described in the methodology and experimental setup sections.

To further assure reproducibility, the implementation is made available in a GitHub repository. The repository contains the code used to generate degraded RSS measurements, apply the compensation method, and evaluate the positioning error. However, reproducibility also depends on the availability of the original DenseVLC dataset and the exact software environment used during the experiments.

C. Use of AI Tools

AI tools were used during the writing process to support grammar correction, phrasing, and structural improvements. The technical content, experimental design, implementation, and interpretation of results remain the responsibility of the author. AI-generated suggestions were reviewed and edited before being included in the thesis.

AI tools were not used as a source of experimental results. All figures and measurements reported in the results section were obtained from the implemented simulation and embedded experiments.

VIII. Discussion

This section discusses the results presented in the results section. The focus is on a correct interpretation of the behavior of the evaluated methods, identifying the trade-offs of the proposed method, and discussing the limitations of this thesis.

A. Interpretation of the Results

The results indicate that compensation designed primarily for gradual aging is vulnerable to sudden or combined degradation. The proposed anomaly correction reduces this vulnerability, but replacement with expected RSS values can increase baseline error. Its principal benefit is therefore stability rather than minimum instantaneous error.

B. Effect of the Right Positioning Model

The results show that the effect of the right positioning model is still important while using the proposed compensation method. The MLP-based implementation remains more stable, while the Random Forest implementation shows larger fluctuations. However, especially at the start while the degradation effects are less prominent, Random Forest can achieve lower errors.

This shows that long-term robustness is not independent of the model chosen alongside the Anti-degradation algorithm. Therefore, robustness under degradation should be considered when selecting a positioning model for long-term VLP deployment

C. Embedded Deployment Trade-Offs

The Raspberry Pi Pico implementation shows that the proposed method can be deployed on constrained hardware. However, the embedded version achieves lower positioning accuracy than the desktop version. This is likely caused by the use of int8 quantization and the reduced measurement pool used for KNN-based expected RSS estimation. Future implementations could investigate whether better quantization methods, a more efficient reference point selection strategy, or adaptive compression can reduce this accuracy loss.

Although quantization and the reduced KNN reference pool decrease positioning accuracy, they allow the embedded calculation to be completed in $3.606\,\mathrm{ms}$ on average. As reported in Table I, the 95th-percentile calculation time is $3.657\,\mathrm{ms}$, while the maximum observed calculation time is $4.008\,\mathrm{ms}$. This demonstrates the trade-off between positioning accuracy, memory usage, and computational latency on resource-constrained hardware.

D. Limitations

The main limitation of this study is the simulation the degradation effect are evaluated through. Although the degradation models are based on mechanisms identified in the literature, real-world degradation may behave differently. In actual deployments, LED aging, contamination, thermal effects, blockages, and failures may interact in different ways than the models used in this thesis.

Another limitation is that the degradation parameters are approximations. The exact rate of LED aging, dust accumulation, thermal droop, and failure behavior depends on the specific hardware, environment, and maintenance conditions. Therefore, the numerical results should be interpreted as an evaluation of the proposed method under simulated degradation scenarios, rather than as exact predictions for every real deployment. However, some of the parameters required for these degradation models, especially those related to LED aging and lumen maintenance, are commonly measured by LED manufacturers. This means that the proposed simulation framework could be adapted to specific LED hardware when such manufacturer data is available. Other parameters, such as dust accumulation, blockages, and maintenance events, remain more deployment-specific and would need to be measured or estimated for each environment.

Finally, the embedded implementation was evaluated on a limited version of the algorithm. The reduced precision due to quantization and a smaller KNN reference pool make the implementation suitable for the Raspberry Pi Pico, but also reduce accuracy. This means that the desktop and embedded results are not directly equivalent.

IX. Conclusion and Future Work

This thesis investigated how structured illumination drift in RSS-based Visible Light Positioning systems can be modeled and compensated for using lightweight algorithms suitable for microcontrollers. Gradual drift was modeled using per-LED scaling factors for aging, optical contamination, and thermal droop, while blockages, failures, and permanent damage were modeled as sudden signal reductions. Gradual drift was compensated for through scaling-based correction, while sudden drift was handled using KNN-based anomaly detection and replacement with expected RSS values.

The method was implemented on a Raspberry Pi Pico using quantized measurements and a reduced KNN reference pool. Its average calculation time was 3.606 ms, demonstrating that lightweight embedded deployment is feasible. Under the combined 100,000-hour degradation scenario, the proposed method produced more stable positioning performance and reduced large error peaks. However, this robustness sometimes resulted in a higher baseline error, while quantization and memory constraints reduced the accuracy of the embedded implementation.

Therefore, structured illumination drift can be modeled as gradual, temporary, and sudden changes in RSS measurements and compensated for on resource-constrained microcontrollers by combining model-based scaling with lightweight anomaly correction. The results demonstrate a trade-off between positioning accuracy, long-term robustness, and embedded resource usage

Future work should evaluate the proposed method in a real long-term VLP deployment. This would make it possible to compare the simulated degradation models with real degradation behavior. Future research could also improve the anomaly detection method, investigate adaptive thresholds, and develop more efficient embedded implementations. In addition, the method could be tested with other positioning models and larger fingerprint databases to evaluate whether the observed trade-offs remain consistent across different VLP setups. Lastly, future work could focus on the combination of our method with the online method of Li[5] to achieve high accuracies without the spikes in positioning error.

References

[1] Jona Beysens et al. "DenseVLC: A cell-free massive MIMO system with distributed LEDs". In: *Proceedings of the 14th International Conference on Emerging Networking EXperiments and Technologies*. 2018, pp. 320–332.

[2] Jiajie Fan, Kam-Chuen Yung, and Michael Pecht. "Predicting long-term lumen maintenance life of LED light sources using a particle filter-based prognostic approach". In: *Expert systems with applications* 42.5 (2015), pp. 2411–2420.

[3] B Hamon and WD Van Driel. "LED degradation: From component to system". In: *Microelectronics Reliability* 64 (2016), pp. 599–604.

[4] SY Ron Hui and YX Qin. "A General Photo-Electro-Thermal Theory for Light Emitting Diode (LED) Systems". In: *IEEE TRANSACTIONS ON POWER ELECTRONICS* 24.8 (2009), p. 1967.

[5] Joey Wenyi Li, Ran Zhu, and Qing Wang. "TinyML-Empowered Indoor Positioning with Light: A Study on the Impact of LED Aging and Failure". TU Delft, 2026.

[6] Junhai Luo, Liying Fan, and Husheng Li. "Indoor Positioning Systems Based on Visible Light Communication: State of the Art". In: ().

[7] Matteo Meneghini et al. "Thermal droop in III-nitride based light-emitting diodes: Physical origin and perspectives". In: *Journal of applied Physics* 127.21 (2020).

[8] Ivan Moreno. "Dirty LED: effect of dust, fat, fingerprints, water, oil and coal on light output". In: *Light-Emitting Diodes: Materials, Devices, and Applications for Solid State Lighting XIV*. Vol. 7617. SPIE. 2010, pp. 234–240.

[9] F Ortenzio, KG Ruddick, and Q Vanhellemont. "Fouling of outdoor optical sensors—the fore-optics contamination experiment". In: *Journal of Optics* 27.12 (2025), p. 125404.

[10] W.D. van Driel et al. "Lumen maintenance predictions for LED packages". In: *Microelectronics Reliability* 62 (2016), pp. 39–44. ISSN: 0026-2714. DOI: https://doi.org/10.1016/j.microrel.2016.03.018. URL: https://www.sciencedirect.com/science/article/pii/S002627141630052X.

[11] Othman Isam Younus et al. "The impact of blocking and shadowing on the indoor visible light positioning system". In: *2021 IEEE 32nd Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*. IEEE. 2021, pp. 1–6.