Indonesian Journal of Data and Science

Volume 7 Issue 2
ISSN 2715-9936
https://doi.org/10.56705/ijodas.v7i2.456

Research Article

An IoT-Based Precision Hydroponic Monitoring System and Long-Term Characterization of Low-Cost Temperature Sensor Drift

Dedy Atmajaya [1,*]; Abdullah Basalamah [2]; Nia Kurniati [3]; Muhammad Iqbal [4]; Thalita Sherly Putri Jasmin [5]

[1] Universitas Muslim Indonesia, Makassar, Indonesia, dedy.atmajaya@umi.ac.id
[2] Universitas Muslim Indonesia, Makassar, Indonesia, abdullah.basalamah@umi.ac.id
[3] Universitas Muslim Indonesia, Makassar, Indonesia, nia.kurniati@umi.ac.id
[4] Universitas Muslim Indonesia, Makassar, Indonesia, 03320230006@umi.ac.id
[5] Universitas Muslim Indonesia, Makassar, Indonesia, 13020230096@umi.ac.id

Correspondence should be addressed to Dedy Atmajaya; dedy.atmajaya@umi.ac.id

Received 10 March 2026; Accepted 30 July 2026; Published 31 July 2026

Abstract:
Introduction: Temperature monitoring is critical in hydroponic cultivation because it influences nutrient solubility, dissolved oxygen, and root uptake, yet low-cost digital sensors commonly used in Internet of Things (IoT) systems may experience accuracy degradation during long-term deployment. This study develops a low-cost IoT monitoring platform for nutrient film technique (NFT) hydroponics and characterizes temperature-sensor drift under continuous operating conditions. **Method**: The system employed two redundant temperature sensors in the nutrient channel and one ambient sensor, with measurements timestamped, filtered, stored locally at the edge, and visualized through a cloud dashboard. Temperature data were recorded hourly for 40 days, producing 960 observations per sensor. Drift was evaluated from the deviation between the primary and redundant channel sensors using mean absolute error, maximum absolute deviation, standard deviation, drift onset, and estimated drift rate. **Results and Discussion**: The primary sensor showed a mean absolute deviation of 0.82 °C over the full monitoring period and a maximum deviation of 1.47 °C. Sensor agreement remained close during the first 10 days but diverged after approximately day 12, with late-period MAE increasing to 1.18 °C and an estimated drift rate of about 0.04 °C/day. **Conclusion**: Long-term drift in low-cost temperature sensors can materially affect hydroponic monitoring accuracy, and redundant sensing provides a practical baseline for future adaptive edge-based calibration methods.

Keywords: Data Quality; Internet of Things; Precision Hydroponics; Sensor Drift; Temperature Sensing

1. Introduction

Networked sensing has moved from a research idea to something growers actually use. The Internet of Things (IoT) now underpins data collection in settings as different as city infrastructure, clinics, and farms [1], [2], and in agriculture it has become a practical way to cut waste and take routine decisions off the operator's hands [3], [4]. Hydroponics is a natural fit for it. The plants grow without soil and live entirely on a managed nutrient solution, so the grower has to watch the environment closely, and a nutrient film technique (NFT) channel, which is little more than a thin, constantly moving film of solution, leaves almost no margin for a bad reading. Sensing here is a requirement, not a convenience.

Of the quantities worth measuring, temperature is among the most consequential. It changes how much nutrient goes into solution, how much oxygen the water carries, and how quickly roots metabolize, so a reading that is a degree off can quietly steer a controller the wrong way. Cost is what usually decides the hardware, and it pushes most builders toward inexpensive digital parts, DS18B20 probes or DHT22 modules wired to an ESP-class microcontroller [5], [6].

Indonesian groups have published a string of such systems recently, from fuzzy-controlled nutrient monitors [7], [8] to rigs aimed at farms in remote areas [9], and our own earlier work paired soil-moisture and air-temperature sensing with a fuzzy controller for shallot irrigation [10]. The same pattern shows up well outside hydroponics, in livestock and in aquaculture monitoring [11], [12].

Cheap sensors share a familiar weakness: they drift. Noise accumulates, the zero point wanders, and accuracy slips over time, and the warm, chemically active air above a nutrient channel only makes matters worse. Calibrating once at the factory, or once before installation, does nothing for degradation that surfaces weeks later. The literature is not short of remedies. Machine learning has spread across agricultural tasks in general [13], learned regressors clearly beat a plain linear fit for low-cost temperature sensors [14], deep models can recover a drift-free signal without any reference at all [15], and in-situ correction has been surveyed specifically for sensor networks [16]. Almost all of these methods, though, begin by assuming the drift is already understood, and many need a central server or a trusted reference to run. Reviews of low-cost sensing keep circling back to one point: stability has to be measured in the field rather than taken on trust [17], [18].

Data quality has itself grown into a research topic within IoT [19], [20], with drift and bias sitting near the top of the problems these frameworks try to catch, and with heavier tools such as reinforcement learning now proposed for cleaning sensor streams [21]. Yet for all that attention, one thing is still missing. Long, multi-sensor temperature records taken from a working hydroponic system, paired with a careful account of how the sensors drift, are hard to find. Most published temperature datasets are short, gathered on a bench, or drawn from urban air-quality networks. Without a common reference of this kind, it is difficult to say whether a new calibration method is really an improvement.

This paper sets out to fill that gap, and it deliberately keeps the spotlight on the sensing system and the data rather than on any one correction algorithm. We make three contributions. First, we describe a low-cost, reproducible IoT monitoring setup for NFT hydroponics that builds in sensor redundancy. Second, we assemble a 40-day, hourly, multi-sensor temperature dataset of 960 samples per sensor, collected while the system was actually running. Third, we characterize the drift of the low-cost sensors, reporting how the bias grows, how the short-term scatter widens, and when the drift sets in, so that later calibration work has a fixed baseline to compete against.

2. Method

The rest of this section walks through the hardware, how readings were collected and cleaned, how the dataset was put together, and the procedure we used to pull the drift signal out of it. Throughout, the design leans toward what survives in the field, where steady power, spare memory, and a reliable connection cannot be assumed.

System Architecture

We organized the system into three layers, with the sensors at the bottom, an embedded board in the middle, and a cloud dashboard on top. The board does the real work. It reads the sensors, stamps each sample with a time, filters it, and writes it to local storage, while the dashboard only shows and archives what the board has already produced. We kept the split deliberate: if the network drops, logging keeps going, because doing the processing close to where the data are born is the whole point of edge computing [22], [23], [24]. **Figure 1** lays out the three layers.

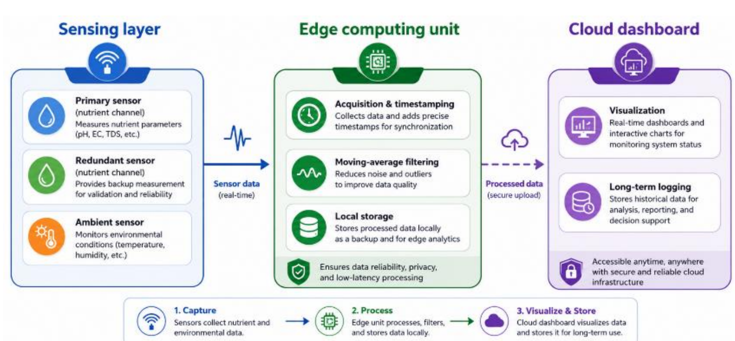

Figure 1. Three-layer architecture of the IoT-based precision hydroponic monitoring platform

Hardware and Sensor Deployment

The temperature sensors are the inexpensive digital type most builders reach for, the same class used in the systems cited above. Instead of adding a laboratory-grade reference, we leaned on redundancy: two identical sensors in the nutrient channel, one primary and one backup, and a third in the surrounding air. Two probes a few centimeters apart in the same flowing film should read almost the same value, so any lasting gap between them tells us about the sensors, not the water. The rig is a standard NFT loop, and a microcontroller-class board handled acquisition and filtering inside a tight power and memory budget. **Figure 2** shows where each sensor sits.

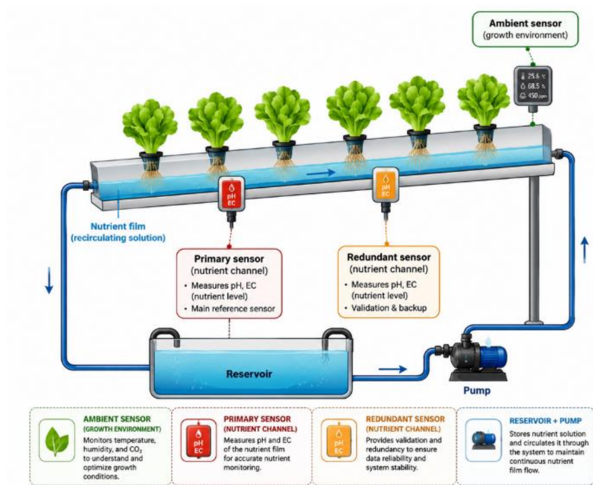

Figure 2. Placement of the primary, redundant, and ambient temperature sensors in the NFT rig

Data Acquisition and Preprocessing

Each sensor was read once an hour and stamped at the board so the series would stay aligned in time. Before storage, a short moving-average window smoothed out spikes and removed obvious outliers; for a window of N samples the filtered value is given by (1).

$$\bar{x}(t) = (1/N)\, \Sigma\, x(t - i), \qquad i = 0 \ldots N - 1 \tag{1}$$

with x(t) the raw reading and x̄(t) the filtered value. The choice of a plain moving average is intentional. Cleaning IoT streams can be done with far heavier machinery [21], but a light filter costs almost nothing on an embedded board and is enough for what we need here. Filtered records went to local storage first and to the dashboard second, and every number reported below comes from the local logs, so nothing depends on whether a given upload happened to succeed.

Data Construction

The system ran without a break for 40 days. At one sample an hour that comes to 960 points per sensor, which is long enough to catch the daily swings, the occasional disturbance, and, most importantly, the slow degradation that a short bench test would never reveal. Each row holds a timestamp, a sensor identifier, the location (channel or air), and both the raw and filtered readings. This is the dataset we analyze, and we plan to release it publicly so the characterization can be reproduced and so later calibration methods have something fixed to be measured on [19], [20].

Drift Characterization Methodology

Two sensors in the same conditions ought to move together; when they stop, one of them is drifting. Writing T_p(t) for the primary reading and T_ref(t) for a reference built from the redundant sensor(s), we take the reference as their median for robustness, as in (2), and define the gap between primary and reference as (3):

$$\bar{x}T\_ref(t) = median\{T\_r1(t), \ldots, T\_rk(t)\} \quad (2)$$

$$d(t) = T_{p(t)} - T_{ref(t)} \quad (3)$$

From d(t) we read three things: its mean over a window, which is the bias; the mean of its absolute value, which is the mean absolute error (MAE) and serves as our accuracy figure; and its rolling standard deviation, which says how jumpy the reading has become. We mark the onset of drift as the first time the gap crosses a threshold delta and stays across it for a window w, as in (4),

$$\bar{x}t * = min\{t : |d(\tau)| > \delta \ for\ all\, \tau \in [t, t + w]\} \quad (4)$$

and, once past that point, we estimate the drift rate from the slope of a straight-line fit to d(t). The persistence window is what keeps a single noisy spike from being counted as real drift.

Evaluation Metrics

We report the MAE, the largest absolute gap, and the standard deviation of d(t) over the whole run, and we split the run into an early stretch (days 1 to 10) and a late one (days 31 to 40) to show how each figure moves. Read together, they say how far the sensor wandered, how erratically, and how fast.

3. Results and Discussion

What follows is the dataset and what the drift analysis turned up. We read the numbers with deployment in mind rather than bench conditions, and because the dataset is being released, anyone can recompute them.

The Dataset at a Glance

Both channel sensors traced the expected daily rhythm, and the air sensor swung wider than the water, which heats and cools more slowly. **Table 1** gives the per-sensor summary over the full 40 days.

Table I. Descriptive statistics of the 40-day temperature dataset, by sensor

Sensor (location)	Mean (°C)	Minimum (°C)	Maximum (°C)	Std. dev. (°C)	Samples
Primary (nutrient channel)	23.8	20.6	27.9	1.86	960
Redundant (nutrient channel)	23.7	20.7	27.6	1.79	960
Ambient (growth environment)	26.4	22.1	31.5	2.53	960

Figure 3 puts a few days of raw and filtered readings side by side. The diurnal shape is plain, and so is the hour-to-hour jitter that the filter takes off.

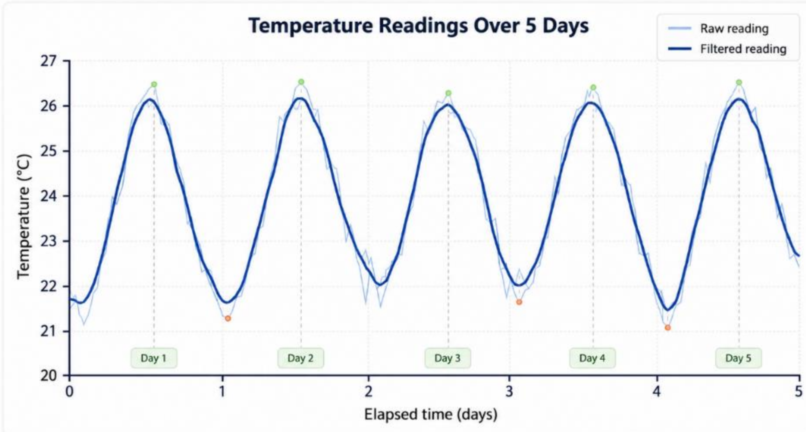

Figure 3. Raw and filtered readings over a representative multi-day window

When the Drift Sets in

For the first week or so the two channel sensors agreed to within roughly a tenth of a degree. Somewhere around day ten to fifteen they began to part, and from there the primary reading pulled steadily away, which is the signature of drift rather than weather and matches what others report for low-cost parts aging in damp conditions [17], [18]. **Figure 4** plots the gap across the whole run; the bend in the curve is the onset.

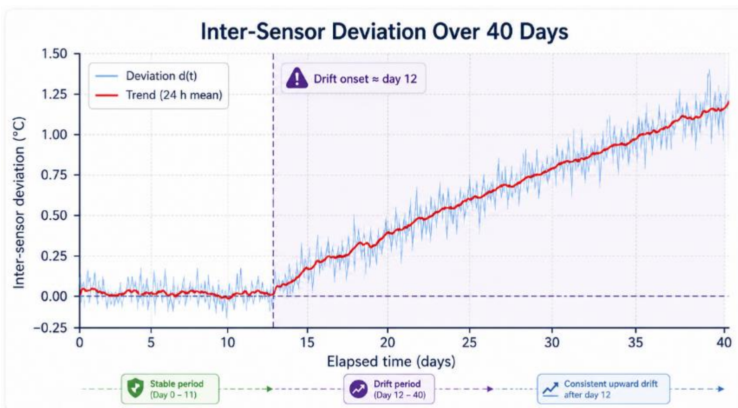

Figure 4. Inter-sensor deviation over the 40-day run; the change in slope marks the drift onset

Putting Numbers on it

Table 2 compares the early and late stretches with the full run. The gap grows by about a factor of ten between the two windows, and both the MAE and the scatter of d(t) climb with it, which is what a real trend looks like rather than passing noise.

Table 2. Drift of the primary sensor relative to the redundant reference

Metric	Early (days 1–10)	Late (days 31–40)	Full 40 days
Mean absolute error, MAE (°C)	0.12	1.18	0.82
Maximum absolute deviation (°C)	0.28	1.47	1.47
Std. dev. of deviation (°C)	0.08	0.33	0.54
Estimated drift rate (°C/day)	—	≈ 0.04	onset ≈ day 12

How this Compares

None of this is out of step with the wider low-cost-sensor literature, which finds accuracy sliding over long deployments and recommends periodic cleaning or recalibration to hold it back [17], [18], [19]. Correction methods are not the missing piece: in-situ schemes have been catalogued for sensor networks [16], learned regressors beat linear correction [14] and deep models recover a usable signal without a reference [15], and Kalman-type filters remain a staple for taming noisy field data [12]. What most of them take for granted is a reference or a central model. Our result comes at the question from the other side. We do not fix the drift here; we measure it on a working system and show its shape, a slow one-way bias with a widening band of noise around it.

For a closed control loop the consequences are practical. A creeping bias slides the effective set-point, so the loop ends up regulating to the wrong temperature, and the growing scatter makes actuators chatter and wastes energy. Because the trouble only appears after the first week or two, a single calibration at install time would never catch it.

Toward Adaptive Calibration

The shape of the drift is encouraging in one respect. It is slow, it shows up clearly in the gap between the two sensors, and it is exactly the sort of thing a static calibration cannot touch. A routine running on the edge board could watch that gap and adjust its correction as the drift appears. Building and testing such a routine is where we are headed next, with the dataset and the baseline here as the yardstick. Releasing the data is meant to let others hold their own methods to the same reference.

4. Conclusion

We built a low-cost IoT monitor for an NFT hydroponic rig and used it to gather and study a 40-day, hourly, multi-sensor temperature record of 960 points per sensor. Reading the primary sensor against its redundant neighbor, we watched it drift away by about 0.82 °C on average and by as much as 1.47 °C, with the gap opening after ten to fifteen days and the short-term scatter widening alongside. The contribution is not another correction algorithm but the groundwork for one: a documented sensing setup, an openly released dataset, and a measured drift baseline that later methods can be tested against. That baseline is the starting point for the adaptive, edge-side calibration we plan to report next.

Acknowledgment

The authors would like to acknowledge the Directorate of Research, Technology, and Community Service (DRTPM), Ministry of Education, Culture, Research, and Technology of the Republic of Indonesia, for providing financial support for this research. The authors also express their sincere appreciation to the Institute for Research and Community Service (LP2S), Universitas Muslim Indonesia, for their support in managing and facilitating this research

References:

[1] L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A survey," Computer Networks, vol. 54, no. 15, pp. 2787–2805, 2010, doi: https://doi.org/10.1016/j.comnet.2010.05.010.

[2] J. Gubbi, R. Buyya, S. Marusic, and M. Palaniswami, "Internet of Things (IoT): A vision, architectural elements, and future directions," Future Generation Computer Systems, vol. 29, no. 7, pp. 1645–1660, 2013, doi: https://doi.org/10.1016/j.future.2013.01.010.

[3] M. S. Farooq, S. Riaz, A. Abid, K. Abid, and M. A. Naeem, "A survey on the role of IoT in agriculture for the implementation of smart farming," IEEE Access, vol. 7, pp. 156237–156271, 2019, doi: https://doi.org/ 10.1109/ACCESS.2019.2949703.

[4] A. L. Duguma and X. Bai, "How the internet of things technology improves agricultural efficiency," Artificial Intelligence Review, vol. 58, Art. no. 63, 2025, doi: https://doi.org/ 10.1007/s10462-024-11046-0.

[5] T. Demir and I. Cicek, "A low-cost and reliable IoT-based NFT hydroponics system using ESP32 and MING stack," Engineering Proceedings, vol. 122, no. 1, Art. no. 3, 2026, doi: https://doi.org/ 10.3390/engproc2026122003.

[6] L. Sitanayah, H. R. M. Joseph, and J. B. Sanger, "A low-cost hydroponic monitoring system with Internet of Things and fuzzy logic," Journal of Computing Theories and Applications, vol. 2, no. 3, pp. 441–455, 2025, doi: https://doi.org/10.62411/jcta.12059.

[7] M. Alimussadad, D. Lestari, and W. M. Utomo, “Smart hydroponic nutrient monitoring and control system using fuzzy logic and IoT,” Jurnal ELTIKOM: Jurnal Teknik Elektro, Teknologi Informasi dan Komputer, vol. 9, no. 2, pp. 115–129, 2025, doi: https://doi.org/10.31961/eltikom.v9i2.1494.

[8] T. A. Farma, D. Palupi, and W. Sirait, “IoT based hydroponic water nutrient and pH monitoring with real-time notifications,” Journal of Innovation Information Technology and Application, vol. 7, no. 2, pp. 340–349, 2025, doi: https://doi.org/10.35970/jinita.v7i2.2940.

[9] F. Laia, F. H. Laia, S. M. Telaumbanua, H. P. Sitorus, and F. J. Hondro, “Development of an IoT-based smart hydroponic monitoring system as a modern agricultural solution in remote areas,” Journal of Intelligent Decision Support System, vol. 8, no. 4, pp. 208–217, 2025, doi: https://doi.org/10.35335/idss.v8i4.318.

[10] A. R. Manga’, D. Atmajaya, and A. Faradibah, “Design and build an automatic spraying system for shallot plants using soil moisture and air temperature sensors with the fuzzy method,” Indonesian Journal of Data and Science, vol. 6, no. 2, pp. 343–355, 2025, doi: https://doi.org/10.56705/ijodas.v6i2.213.

[11] A. F. Shimbun, M. A. Alfian, A. S. Jati, and E. Faizal, “An optimization strategy for reducing CO2 in livestock farming with IoT integration and decision support system approach using linear programming,” Indonesian Journal of Data and Science, vol. 6, no. 1, pp. 46–54, 2025, doi: https://doi.org/10.56705/ijodas.v6i1.204.

[12] M. Misbahuddin, N. Cokrowati, M. S. Iqbal, O. Farobie, A. Amrullah, and L. Ernawati, “Kalman filter-enhanced data aggregation in LoRaWAN-based IoT framework for aquaculture monitoring in Sargassum sp. cultivation,” Computers, vol. 14, no. 4, Art. no. 151, 2025, doi: https://doi.org/10.3390/computers14040151.

[13] A. Kamilaris and F. X. Prenafeta-Boldu, “Deep learning in agriculture: A survey,” Computers and Electronics in Agriculture, vol. 147, pp. 70–90, 2018, doi: https://doi.org/10.1016/j.compag.2018.02.016.

[14] F. Nan, C. Zeng, H. Shen, and L. Lin, “Calibration of integrated low-cost environmental sensors for urban air temperature based on machine learning,” Sensors, vol. 25, no. 11, Art. no. 3398, 2025, doi: https://doi.org/10.3390/s25113398.

[15] Y. Wang, A. Yang, X. Chen, P. Wang, Y. Wang, and H. Yang, “A deep learning approach for blind drift calibration of sensor networks,” IEEE Sensors Journal, vol. 17, no. 11, pp. 3432–3442, 2017, doi: https://doi.org/10.1109/JSEN.2017.2703885.

[16] B. Maag, Z. Zhou, and L. Thiele, “In situ calibration algorithms for environmental sensor networks: A review,” IEEE Sensors Journal, vol. 19, no. 13, pp. 4924–4938, 2019, doi: https://doi.org/10.1109/JSEN.2019.2905910.

[17] J. A. Abdinoor, Z. K. Hashim, B. Horvath, S. Zsebo, D. Stencinger, G. Hegedus, L. Bede, A. Ijaz, and I. M. Kulmany, “Performance of low-cost air temperature sensors and applied calibration techniques — A systematic review,” Atmosphere, vol. 16, no. 7, Art. no. 842, 2025, doi: https://doi.org/10.3390/atmos16070842.

[18] M. Naloufi, T. Abreu, S. Souihi, C. Therial, N. Angelotti de Ponte Rodrigues, A. Guillot Le Goff, M. Saad, B. Vincon-Leite, P. Dubois, M. Delarbre, P. Kennouche, and F. S. Lucas, “Long-term stability of low-cost IoT system for monitoring water quality in urban rivers,” Water, vol. 16, no. 12, Art. no. 1708, 2024, doi: https://doi.org/10.3390/w16121708.

[19] H. Y. Teh, A. W. Kempa-Liehr, and K. I.-K. Wang, “Sensor data quality: a systematic review,” Journal of Big Data, vol. 7, Art. no. 11, 2020, doi: https://doi.org/10.1186/s40537-020-0285-1.

[20] L. Zhang, D. Jeong, and S. Lee, “Data quality management in the Internet of Things,” Sensors, vol. 21, no. 17, Art. no. 5834, 2021, doi: https://doi.org/10.3390/s21175834.

[21] A. F. Y. Mohammed, S. M. Sultan, J. Lee, and S. Lim, “Deep-reinforcement-learning-based IoT sensor data cleaning framework for enhanced data analytics,” Sensors, vol. 23, no. 4, Art. no. 1791, 2023, doi: https://doi.org/10.3390/s23041791.

[22] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, “Edge computing: Vision and challenges,” IEEE Internet of Things Journal, vol. 3, no. 5, pp. 637–646, 2016, doi: https://doi.org/10.1109/JIOT.2016.2579198.

[23] M. Satyanarayanan, “The emergence of edge computing,” Computer, vol. 50, no. 1, pp. 30–39, 2017, doi: 10.1109/MC.2017.9.

[24] W. Shi and S. Dustdar, “The promise of edge computing,” Computer, vol. 49, no. 5, pp. 78–81, 2016, doi: https://doi.org/10.1109/MC.2016.145.