IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 19 June 2024, accepted 12 July 2024, date of publication 16 July 2024, date of current version 25 July 2024.

Digital Object Identifier 10.1109/ACCESS.2024.3429197

RESEARCH ARTICLE

Toward a Performance-Based Trustworthy Edge-Cloud Continuum

INDIKA DHANAPALA, (Member, IEEE), SOURABH BHARTI, (Member, IEEE), ALAN MCGIBNEY, (Member, IEEE), AND SUSAN REA, (Senior Member, IEEE)
Nimbus Research Centre, Munster Technological University, Cork, T12 P928 Ireland

Corresponding author: Indika Dhanapala (Indika.Dhanapala@mtu.ie)

This work was supported in part by the Science Foundation Ireland (SFI) as part of the SFI Research Centre for Future Networks and Communications (CONNECT) under Grant 13/RC/2077_P2 and in part by the European Union's Horizon 2020 Research and Innovation Programme under Grant 958339.

ABSTRACT The Edge-Cloud Continuum refers to the dynamic provisioning of distributed computing and network resources that can be scaled to support the creation of secure, resource-efficient, and decentralised digital ecosystems, which also support federated topologies for collaborating and sharing resources. Trusted interaction and orchestration of distributed edge-cloud resources are the fundamental principles of distributed network infrastructure and service provisioning. The zero trust architecture (ZTA) paradigm is gaining momentum based on being able to ensure trusted and secure interaction for edge-cloud networks. However, ZTA's strict authentication policy mandates devices to be authenticated for every session, leading to significant overhead for resource-constrained devices engaged in multiple sessions. To address this challenge, this paper proposes a ZTA that integrates a performance-based trust assessment mechanism, allowing a higher number of consecutive sessions without the need for costly authentication/authorisation while preserving system integrity. Reputation, viewed through the performance lens, is a metric to gauge a node's trustworthiness, considering its past behaviour and interactions. The proposed trust assessment mechanism is evaluated for its feasibility within our conceptualised ZTA for edge computing environments with limited resources, and simulation results demonstrate the practicality of utilising this technique in zero trust environments.

INDEX TERMS Edge-cloud continuum, edge computing, reputation, performance-based trust, zero trust architecture.

I. INTRODUCTION

Advancements in technology and the rise of data-intensive applications have spurred the evolution of computing architecture. Traditional approaches such as centralised cloud computing and localised edge computing fall short of meeting the diverse demands of modern applications. The edge-cloud continuum has emerged as a new paradigm, blending edge computing and cloud computing resources into a cohesive system. This integration effectively closes the gap between localised data processing and the centralised computing infrastructure. It significantly boosts the efficiency of distributed infrastructure and service delivery while also facilitating the creation of decentralised and federated network topologies.

Conventional perimeter-based security mechanisms assume that all devices and users connected to a local network are trusted [1] and can be granted access to resources within the network and external access to services provided in the local network based on verification using security and identification mechanisms. However, these assumptions have become increasingly inaccurate as cyberattacks have become more sophisticated and persistent [2], [3], [4]. Perimeter-based security is becoming less effective as *i*) protecting the network boundary is no longer sufficient to protect data and applications, *ii*) with an increasing number of mobile and remote devices, network perimeters are becoming less defined, and *iii*) irrespective of being intentional or not, insider threats are a growing concern.

Simultaneously, the rise of the zero trust security model has revolutionised how organisations approach network security. This model emphasises strict access controls and

The associate editor coordinating the review of this manuscript and approving it for publication was Hadi Tabatabaee Malazi.

continuous verification of devices and users [5], [6], [7], [8]. By combining the strengths of the edge-cloud continuum and the zero trust security model, organisations can achieve a resilient and secure computing environment. This integration enables localised processing and data analysis at the edge while leveraging the cloud's scalability, storage capacity, and other advanced capabilities.

Zero trust is an alternative security approach that assumes that no devices or users, including those within a local network, can be trusted [2], [4]. Therefore, all devices and users, regardless of their location (physical or virtual), must be authenticated and authorised before being granted access to resources [9]. In essence, zero trust operates on the premise that trust is never implicitly granted but must be continuously evaluated to ensure resource protection [9]. The key principles of a ZTA include [3]:

- Identity verification – All users and devices must undergo authentication and authorisation before accessing resources.
- Least privilege access – Users and devices should only be given the minimum required access to perform their tasks.
- Microsegmentation – To limit the attack surface, the network is divided into small isolated segments with strict access controls between them.
- Continuous monitoring – All activities in the network are continuously monitored for any sign of suspicious activity.
- dynamic policy enforcement – Resource access policies should be continuously updated to align with changes in a user's or device's identity or behaviour.

Due to the continuous monitoring of the entities interacting with the system, ZTA requires more computational resources compared to the traditional perimeter-based approaches.

Reputation management, serving as a trust assessment mechanism, finds widespread application across various domains, including e-commerce and social networks. It involves gathering and analysing feedback and opinions from other users/devices regarding a particular entity, such as a person, product, or service, and generating a score based on this feedback. Using the same principles, reputation management can also be used to evaluate the trustworthiness of nodes in a distributed network, where each node is assigned a reputation score based on its past behaviour and interactions with other nodes [10]. Nodes with high reputation scores are regarded as more trustworthy, while those with low scores are considered less trustworthy. In addition, reputation management can also be used in conjunction with other trust assessment mechanisms, such as context / behaviour-based approaches, to provide a more comprehensive trust evaluation. By leveraging feedback and opinions from other nodes, reputation management provides a more accurate and dynamic assessment of trustworthiness, which can help improve the security and reliability of the distributed network.

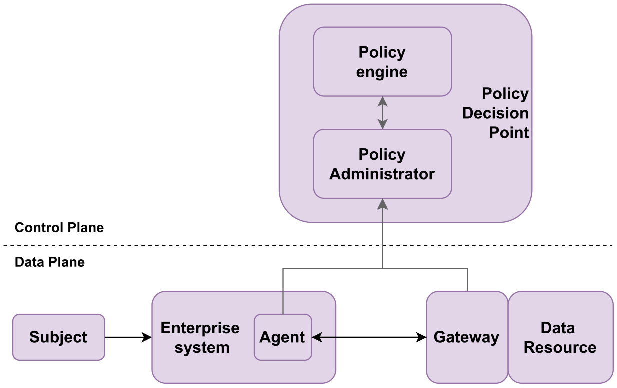

FIGURE 1. NIST zero trust architecture [9].

A. BACKGROUND AND CONTEXT

1) OVERVIEW OF NIST ZERO TRUST ARCHITECTURE

The NIST ZTA focuses on authentication, authorisation, and reducing implicit trust zones to prevent unauthorised access to data and services while enforcing the least privileged access [9]. All of these tasks must be performed while maintaining availability and minimising delays in authentication mechanisms. To achieve this, NIST proposes to have a policy decision point (PDP) and a policy enforcement point (PEP) to manage access. The former is responsible for making decisions to grant access to data/service for a client, while the latter enables, monitors and terminates the connections between clients and services. Although there are various deployment options for the PEP, this article follows the device agent/gateway-based deployment model (depicted in Figure 1) as it is best suited for discrete resources capable of communicating with a gateway. This model divides the PEP into two components: the agent and the gateway. The agent, a software component installed on devices requesting access to data/services, directs traffic to PEP for evaluation and coordinates connections. The gateway communicates with the PDP and allows only approved connections as determined by the PDP. In this manner, the agent and the gateway collectively function as the PEP.

2) MODIFIED ZERO TRUST FRAMEWORK

This work presents a *modified zero trust architecture* as shown in Figure 2 that improves the existing NIST ZTA [9] (see Figure 1) with a performance-based and ML-based mechanism. It facilitates trust assessment for managing and orchestrating next-generation digital services across the edge-cloud continuum.

The proposed architecture introduces performance-based reputations for nodes to enable multiple consecutive sessions in a zero trust network. The number of sessions allowed depends on the reputation score of the requesting node. The underlying motivation for utilising reputation and contextual data in a zero trust network is founded on the hypothesis that *'by leveraging a score and contextual-based approach,*

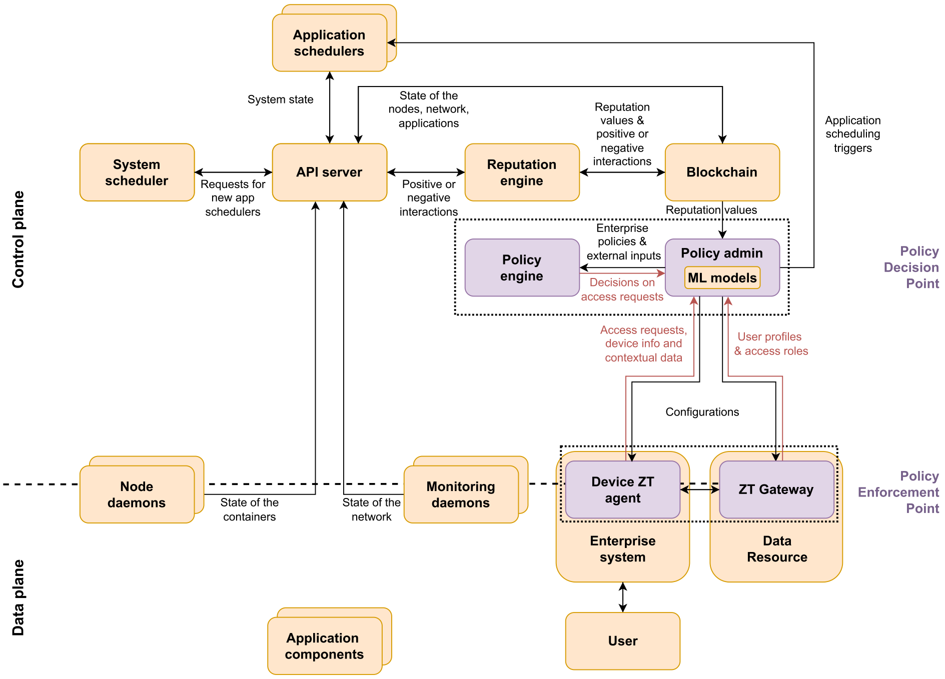

FIGURE 2. Modified zero trust architecture.

a performance-based trust assessment can provide more dynamic and granular access control in a zero trust network, which can in turn increase the number of consecutive sessions per node before re-evaluating its trustworthiness'.

Implementing a trust evaluation algorithm that combines contextual and score-based attributes enables dynamic and fine-grained access control. The scoring mechanism provides a confidence level for the requesting node and adapts quickly to changing factors, unlike static access policies [7], [9]. Moreover, contextual data allow the detection of anomalous behaviour in nodes and the network, enhancing the security of active sessions.

The reputation engine evaluates node reputation based on their positive and negative interactions, collected by monitoring daemons (Figure 2). Positive interactions in the ML context can include meeting expected model accuracy during training or obtaining predictions within a specific delay. Reputation information is shared through a Blockchain, allowing distributed computation of indirect node reputation (see Section III for details).

The proposed modified ZTA adopts the NIST zero trust framework's device agent/gateway model due to its ability to safeguard borderless networks and seamless integration with the edge-cloud architecture [9]. Software agents, namely *device zero trust agents* and *monitoring daemons* (Figure 2), collect contextual and behavioural data for individual nodes, aiding the policy administrator in decision-making. ML models trained with contextual data detect abnormal node behaviour. ML models and reputation scores help the policy engine to make access decisions, dynamically adapt policies during active sessions, and improve security. If access is granted, the application schedulers schedule components for service provisioning. In addition, the policy administrator reconfigures the policy enforcement point to reflect the current trustworthiness of the device. Furthermore, the system scheduler creates a new application scheduler when receiving a new application request. The Blockchain tracks the network, nodes, and running applications' status, offering deeper insights into the system's functionality. The logical components in Figure 2 are containerised and deployed within a single edge server, while multiple of these edge servers are distributed across the edge-cloud continuum. Note that Figure 2 illustrates the key elements of the full ZTA that has been proposed. This paper specifically focuses on the policy decision point (PDP), where the policy engine utilises the reputation scores to inform its decision-making process.

The importance of reputation in the proposed architecture lies in its role of determining the number of active sessions granted to a trustworthy node before reevaluating its trustworthiness. Inaccurate reputations could compromise active sessions, posing a significant threat to the entire network. Hence, evaluating the performance of the reputation engine is crucial and is the focus of this paper.

B. OUR CONTRIBUTIONS

While the article delves into conceptual aspects, its foundation lies in the addressing of practical challenges in real-world IoT applications. Coordinating IoT devices, edge servers, and cloud servers, often operated by different vendors, is complex. The integration of Blockchain, as depicted in Figure 2, ensures transparent and secure sharing of reputation-related data among nodes, regardless of vendor differences. Our framework establishes trust dynamically through its performance-based system, enabling efficient and secure interactions, even in environments where devices and servers are managed by different entities. Recognising the importance of practical application scenarios, future work will involve testing and refining our framework in real-world IoT environments to validate its effectiveness and adaptability to diverse vendor landscapes.

The contributions of this paper are as follows:

- An algorithm for performance-based trust assessment, expanding on the model introduced in [10] by incorporating interaction frequency and contextual factors.
- Analysis of the complexity of the proposed trust assessment algorithm.
- Evaluation of the performance of the trust assessment algorithm in a realistic simulation scenario.

The following sections of this paper include a discussion of related work in Section II, the details of reputation calculation in Section III, and simulation scenarios and performance metrics in Section IV. The feasibility of the reputation model in the modified ZTA is evaluated in Section V, followed by conclusions and future research directions in Section VI.

II. REPUTATION MANAGEMENT

Numerous recent studies have proposed solutions to tackle distributed reputation management in the cloud-edge continuum. Examples of such works encompass research conducted by Yuan and Li [11], DREAMS [12], Liang et al. [13], COMITMENT [14], Latif et al. [15], Liu et al. [16], Feng et al. [17], RTEM [18], TOF [19], ETARR [20], Guo et al. [21], TALMSC [22], and Kang et al. [10]. These studies are detailed below.

Yuan and Li [11] presents a novel approach to distributed trust computing by leveraging multi-source feedback from diverse devices in edge computing. Unlike our approach, which utilises subjective logic and can handle uncertainty, the authors employ an objective information entropy theory-based fusion algorithm to compute trust values. Their trust computation incorporates direct trust and recommendation, considering historical interactions and evaluating QoS requirements at the network edge. Their mechanism performs well in countering bad-mouthing attacks originating from malicious feedback providers. DREAMS [12] proposes a distributed reputation management scheme for vehicular edge computing. In this scheme, reputation calculation is performed using subjective logic. Each reputation score received from other devices is assigned a weight based on factors such as familiarity, similarity, and timeliness. These factors determine the level of prior knowledge that the rater has about the ratee, consider the conditions of the rater, and take into account the freshness of the reputation information. In [13], a trust computing mechanism is introduced to minimise trust computing overhead, communication overhead, and communication delays. This is achieved through the integration of multi-source feedback and fog computing. The direct trust computation in this approach relies on service quality, with node preferences serving as weights. Additionally, the recommendations are calculated on the basis of historical data. The final trust score is obtained through a weighted summation of the direct trust and recommendation. COMITMENT [14] presents a distributed trust and recommendation model for fog computing. This model incorporates both direct and indirect periodic reputation calculations derived from previous interactions with other nodes. The final reputation score is determined by a weighted average of recommendations provided by peers, the weight assigned based on the level of trustworthiness of each peer. Reference [15] presents a trust management model for edge computing that comprises a rating management module and a trust calculation module. Node ratings are determined by evaluating quality of service (QoS) parameters, and the trust score is calculated as a weighted average of these parameters. RTEM [18] proposes a three-tier trust evaluation framework for mobile edge computing networks, wherein three types of trust evaluation (i.e., identity, capability, and behaviour) are carried out. In this approach, firstly, the authenticity of the node is confirmed; secondly, the capability of the device to run tasks is determined; lastly, the behaviour trust is calculated based on historical interactions and local/global reputation. Their trust evaluation model considers direct trust, indirect trust, and rater credibility. In TOF [19], [23], trustworthiness is used to cluster IoT smart objects (SOs) to introduce a competitive environment and to promote correct behaviours. To quantify trustworthiness, reputation is used and is calculated in edge servers based on historical feedback received from the respective counterparts of individual SOs. Their reputation model considers the reliability of providing honest feedback and collusive behaviour of SOs, enabling TOF to tackle performance against malicious activities. To improve the efficiency of emergency task completions and reduce the associated maintenance costs, ETARR [20] proposes a reputation-based mechanism for allocating emergency tasks to maintenance personnel. Work enthusiasm and work activities are used as indicators in their reputation model and employ an LSTM model to predict the reputation value. Reference [21] proposes a trust management model for cloud environments based on mutual trust and uses a reward-with-punishment mechanism to eliminate malicious entities. Direct trust, recommendation trust, rater credibility, comprehensive trust, and trust trend are explored to calculate mutual trust values in their trust model. TALMSC [22] proposes a trust-based agent learning model that uses fuzzy comprehensive evaluation (FCE) techniques to determine trust values. Their trust model considers trust features such as subjectivity, objectivity, uncertainty, context sensitivity, and QoS factors. In [10], subjective logic is applied to select a set of top-performing workers for reliable federated learning. The reputation management system is decentralised and implemented using a Blockchain.

The reputation management scheme for ZTA used in our article builds upon this work.

All the aforementioned works propose solutions for distributed trust management. However, none of these works is specifically designed for use in zero trust environments, nor have they employed zero trust principles. Zero trust environments require characteristics such as transparency, interoperability, scalability, privacy preservation, resilience to attacks, and context awareness, which are not specifically addressed in these works. SeComTrust [24] proposes a trust management model based on zero trust principles in a community cloud, where organizations interact with each other to share resources. In this approach, the community cloud is subdivided into three groups based on the sensitivity levels of their resources. SeComTrust assesses trust relationships between the organizations through subjective logic and provides a mechanism to update trust values, allowing promotion or relegation within the three groups. To establish secure information sharing, [16] introduces a zero trust mechanism that uses smart contracts as a voting system and a subsequent computation of node reputation. The node with the highest reputation is included in the reputation chain, and these values are utilised for future transactions. Furthermore, data verification via the Blockchain is employed to penalise malicious nodes and ensure the integrity of the system. Nonetheless, the voting mechanism is an extra process which introduces an additional overhead. On the contrary, [25] proposes a task offload mechanism for vehicular services based on zero trust principles, where the reputation of roadside units (RSUs) is determined using subjective logic, taking into account the analysis of historical service records in a Blockchain. This approach does not require an additional mechanism for data collection to calculate reputation. Instead, it uses existing historical service records. Similarly, [17] introduces a reputation evaluation approach to reduce the possibility that nodes access malicious cloud services that take into account various factors such as computational capability, storage capability, service capability, and network performance of the cloud service. The reputation of cloud services is determined as a weighted summation of these dimensions, considering the quality of service (QoS) requirements. In addition, the reputation evaluation mechanism incorporates historical reputation values, where the results obtained in previous test cycles are weighted and combined with the current reputation value to obtain the combined reputation value. The proposed trust management model is applied within a ZTA in conjunction with a Blockchain, which serves as distributed storage for user identities.

Similar to the works in [17] and [25], the modified ZTA proposed in this paper employs a Blockchain for disseminating reputation and other pertinent data across the network, ensuring transparency and security of reputation information. In addition, our approach incorporates additional features to enhance the accuracy of the reputation calculation.

The security challenges associated with trust management systems (TMS) have a broad impact on the entire system. Therefore, it is crucial to employ techniques to combat attacks from malicious devices that seek to spread false trust reports. Among the common trust-based attacks, self-promoting attacks, bad-mouthing attacks, ballot stuffing attacks, and deceptive misrepresentation attacks are prevalent [26], [27]. A self-promoting attack occurs when a malicious device attempts to illicitly enhance its reputation by providing dishonest recommendations. In bad-mouthing attacks, attackers intentionally give negative ratings to a benign node to diminish its trust score. Inversely, a ballot-stuffing attack involves manipulating the reputation of compromised devices to elevate their trustworthiness levels. Deceptive misrepresentation involves a malicious node deliberately providing false information about its capabilities to manipulate the reputation system.

There are several recent works focused on mitigating trust attacks in TMSs. For example, Trust2Vec [28] can handle large-scale bad-mouthing and self-promoting attacks by leveraging a random-walk network exploration algorithm. It serves as a robust safeguard for network security and data integrity, acting as a referee that promotes trustworthy devices while punishing any malicious activities. Reference [27] presents a threat model to detect bad-mouthing and ballot-stuffing attacks with the help of distributed hash tables.

Accurate reputation calculation and effective defence against trust-based attacks are prerequisites for dynamic and granular access control in the modified ZTA.

III. PERFORMANCE-BASED TRUST ASSESSMENT

Reputation in communication networks measures a node's trustworthiness based on its past behaviour and interactions, reflecting how effectively it fulfilled its duties and responsibilities within the network. Applying this principle in the edge-cloud continuum: *i*) end devices can use edge server reputation to select suitable servers for task offloading. *ii*) edge servers can utilize end device reputation to decide on granting access to their resources. Thus, each node in the edge-cloud continuum must compute the reputation of the nodes it intends to interact with.

After interacting with other nodes on the network for the execution of tasks, each node marks the quality of its interaction. An interaction is marked positive if it meets the expected level of processing delay, communication latency and accuracy, otherwise, it is considered negative. The reputation computation takes into account both positive and negative interactions, denoted by $\alpha_{i\to j}^{t_y}$ and $\beta_{i\to j}^{t_y}$, respectively.

A. DIRECT REPUTATION

The Reputation opinion of node i towards a peer node j within the time window t_y consists of i) belief degree ($b_{i\to j}^{t_y}$) – the probability of node i believing that node j will successfully complete a task; *ii*) disbelief degree ($d_{i\to j}^{t_y}$) – the probability of node i not believing that node j will successfully complete a task; and *iii*) uncertainty ($u^{t_y}_{i\to j}$) – node i's level of uncertainty regarding node j's ability to successfully complete a task [10]. These terms are derived from the definition of an opinion in the subjective logic technique [29]. Equations 1, 2, 3, and 4 establish the connection between the belief degree, disbelief degree and uncertainty with the positive and negative interactions, where $PDR^{t_y}_{i\to j}$ is the packet delivery ratio between node i and j within t_y time window. The belief and disbelief values, respectively, are the percentage of positive and negative interactions between node i and j within the scope of confirmation, $(1-u^{t_y}_{i\to j})$.

$$b^{t_y}_{i\to j} = (1-u^{t_y}_{i\to j})\frac{\alpha^{t_y}_{i\to j}}{(\alpha^{t_y}_{i\to j}+\beta^{t_y}_{i\to j})}, \tag{1}$$

$$d^{t_y}_{i\to j} = (1-u^{t_y}_{i\to j})\frac{\beta^{t_y}_{i\to j}}{(\alpha^{t_y}_{i\to j}+\beta^{t_y}_{i\to j})}, \tag{2}$$

$$u^{t_y}_{i\to j} = 1 - PDR^{t_y}_{i\to j}, \tag{3}$$

$$b^{t_y}_{i\to j} + d^{t_y}_{i\to j} + u^{t_y}_{i\to j} = 1 \tag{4}$$

The expected value of belief degree is considered to be the direct reputation value of node i towards node j within the time frame t_y. Direct reputation (*DR*) is determined by Equation 5, wherein γ $(0 \le \gamma \le 1)$ represents the effect of uncertainty on reputation.

$$DR^{t_y}_{i\to j} = b^{t_y}_{i\to j} + \gamma u^{t_y}_{i\to j} \tag{5}$$

Several factors can affect the reputation of a node, such as how frequently a node interacts with another (*interaction frequency*), how recently a node has interacted with another (*interaction freshness*) and the context in which the interaction occurred (*interaction context*). Nodes that frequently interact with others are more likely to be trustworthy as they have established a track record of successful interactions. Conversely, nodes that rarely interact with others may be less reliable and therefore have a lower reputation. Nodes that have not interacted with others for a long time may have outdated information. As a result, their reputation score may be less reliable than nodes that have more recent interactions. Furthermore, by taking the context of the interactions into account, the reputation model can better understand the nature and purpose of the interactions, resulting in an improvement in its accuracy and effectiveness in assessing the trustworthiness of a node. Therefore, the accuracy of the reputation is directly linked to the interaction frequency and the interaction freshness, with higher values of these factors resulting in more reliable reputation measurements. In addition, the relevance of the interaction context also plays a role in determining accuracy. Consequently, incorporating interaction frequency, freshness, and context can increase the precision of reputation calculations.

Equation 6 calculates the interaction frequency (*IF*) between node i and j within the time window t_y, considering positive and negative interactions. The resulting *IF* is used as a weight during the calculation of the reputation opinions in Equations 1, 2, and 3.

$$IF^{t_y}_{i\to j} = \frac{(\alpha^{t_y}_{i\to j}+\beta^{t_y}_{i\to j})}{\frac{1}{|I|}\Sigma_{s\in I}(\alpha^{t_y}_{i\to s}+\beta^{t_y}_{i\to s})}, \tag{6}$$

where I is the set of all nodes that interacted with node i.

To incorporate interaction freshness in the calculation, a freshness fading function θ_y is used. The function is defined as $\theta_y = Z^{Y-y}$, where $Z \in (0,1)$ and $y \in [1, Y]$. Here, Y represents the total number of time slots considered for the calculation.

$$b_{i\to j} = \frac{\sum_{y=1}^{Y}\theta_y b^{t_y}_{i\to j}}{\sum_{y=1}^{Y}\theta_y} \tag{7}$$

$$d_{i\to j} = \frac{\sum_{y=1}^{Y}\theta_y d^{t_y}_{i\to j}}{\sum_{y=1}^{Y}\theta_y} \tag{8}$$

$$u_{i\to j} = \frac{\sum_{y=1}^{Y}\theta_y u^{t_y}_{i\to j}}{\sum_{y=1}^{Y}\theta_y} \tag{9}$$

$$DR_{i\to j} = \frac{\sum_{y=1}^{Y}\theta_y DR^{t_y}_{i\to j}}{\sum_{y=1}^{Y}\theta_y} \tag{10}$$

The interaction context refers to the nature of interactions between two entities within a network, including model training, model inference, and other (interactions outside these categories). For example, if two nodes collaborate on tasks related to machine learning (ML) model training, their interaction context would be labelled as ML model training. In such a scenario, if node i intends to delegate its new ML model training tasks to node j, node i assigns a weight δ $(0.5 \le \delta \le 1)$ to the computed reputation score of node j only if node j has prior ML model training interactions with node i. Otherwise, a weight of $(1-\delta)$ is used.

B. SIMILARITY OF RECOMMENDERS

There are situations where two nodes within a network may not have engaged in any prior interactions, yet there arises a necessity to assess each other's level of trustworthiness. This is accomplished using indirect reputation (*IR*), as illustrated in Figure 3, where a common node acts as a recommender. There may be multiple recommenders for the same node j, and they may not all have the same level of credibility, so individual weights are assigned to each recommender based on their similarity. Equation 11 computes the similarity

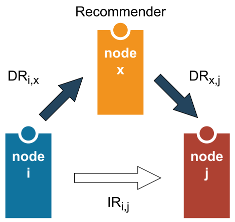

FIGURE 3. Indirect reputation via a recommender.

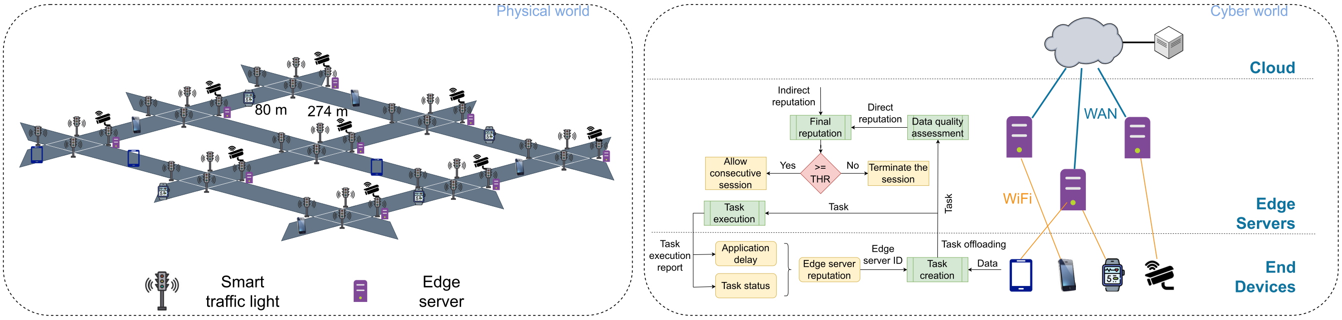

FIGURE 4. Simulation scenario. The left figure represents the physical scenario of the city, while the right figure maps the scenario onto the corresponding layered communication network and its functional diagram.

of a recommender x, where I and X are the sets of nodes that have had interactions with node i and x, respectively, while C represents the set of nodes that have had interactions with both node i and x (i.e., $C = I \cap X$).

$$sim(i, x) = \frac{\Sigma_{k\in C}(DR_{i\to k} - \overline{DR_i})(DR_{x\to k} - \overline{DR_x})}{\sqrt{\Sigma_{k\in I}(DR_{i\to k} - \overline{DR_i})^2}\sqrt{\Sigma_{k\in X}(DR_{x\to k} - \overline{DR_x})^2}} \quad (11)$$

C. INDIRECT REPUTATION

Subjective logic is utilized to compute indirect reputation (IR), and Equations 12, 13, and 14 show how to determine indirect reputation opinions (i.e., reputation opinions given by recommenders) using similarity as a weight. Indirect reputation is then calculated in a manner comparable to Direct Reputation in Equation 5.

$$b^{Rec}_{x\to j} = \frac{1}{\Sigma_{x\in C}Sim(i, x)}\Sigma_{x\in C}Sim(i, x)b_{x\to j} \quad (12)$$

$$d^{Rec}_{x\to j} = \frac{1}{\Sigma_{x\in C}Sim(i, x)}\Sigma_{x\in C}Sim(i, x)d_{x\to j} \quad (13)$$

$$u^{Rec}_{x\to j} = \frac{1}{\Sigma_{x\in C}Sim(i, x)}\Sigma_{x\in C}Sim(i, x)u_{x\to j} \quad (14)$$

D. FINAL REPUTATION

The final reputation opinions are obtained by merging direct and indirect reputation opinions, which is accomplished using Equations 15, 16, and 17. Subsequently, the final reputation value is computed using Equation 18, analogous to Equation 5.

$$b^{final}_{i\to j} = \frac{b_{i\to j}u^{Rec}_{x\to j} + b^{Rec}_{x\to j}u_{i\to j}}{u_{i\to j} + u^{Rec}_{x\to j} - u_{i\to j}u^{Rec}_{x\to j}} \quad (15)$$

$$d^{final}_{i\to j} = \frac{d_{i\to j}u^{Rec}_{x\to j} + d^{Rec}_{x\to j}u_{i\to j}}{u_{i\to j} + u^{Rec}_{x\to j} - u_{i\to j}u^{Rec}_{x\to j}} \quad (16)$$

$$u^{final}_{i\to j} = \frac{u_{i\to j}u^{Rec}_{x\to j}}{u_{i\to j} + u^{Rec}_{x\to j} - u_{i\to j}u^{Rec}_{x\to j}} \quad (17)$$

$$FR^{final}_{i\to j} = b^{final}_{i\to j} + \gamma u^{final}_{i\to j} \quad (18)$$

E. COMPLEXITY ANALYSIS

Devices within a communication network employ reputation calculations before establishing connections with other devices. For instance, an end device utilises reputation calculations when selecting an edge server before offloading tasks. These calculations are executed by devices within the network (refer to Figure 4) in four steps, outlined in Algorithm 1. The complexity involved in each step is analysed to understand the computation delay. Additionally, communication complexity is used to infer communication latency, as discussed in Section III-E2. These two metrics are defining parameters of a trusted edge cloud continuum, which is the ultimate objective of the proposed modified ZTA.

Algorithm 1 Reputation of Node i on Node j

```
**Input**: $i$; $j$; $\gamma$ (Eq. 5); $\delta$ (Eq. 10); $Z$: freshness decay function; $Y$: number of time slots considered; DB: reputation database (Blockchain);
**Output**: $FR^{t_y}_{i\to j}$
**while** true **do**
    // calculate direct reputation between all nodes
    update interactions between nodes within $t_y$
    calculate direct reputations according to Eq. 5
    update DB with the latest reputation figures
    // indirect reputation between a recommender $x$ and $j$
    iterate over $I$ and calculate interaction frequencies of $i$ (Eq. 6)
    iterate over $C, I$ and $X$ and calculate the similarity between $i$ and $x$ (Eq. 11). Add $IF$, freshness, and context into the calculation of the similarity when using $DR$.
    iterate over $C$ and calculate indirect reputation of $x$ and $j$ (Eq. 12– 14 and Eq. 5)
    // final reputation
    $FR^{t_y}_{i\to j}$ ← combine direct and indirect reputation opinions and calculate $FR$ (Eq. 18). Add $IF$, freshness, and context into the calculation of the final reputation when using $DR$.
    return $FR^{t_y}_{i\to j}$
**end while**
```

1) COMPUTATIONAL COMPLEXITY

As outlined in Algorithm 1, the computational complexity of direct reputation predominantly arises from the interaction frequency calculation, which involves iterations over peers. Conversely, both interaction freshness and interaction context have a constant complexity of $O(1)$. Consequently, the computational complexity for direct reputation (DR) computation can be expressed as $O(|I| + |Y|)$. Here, $|I|$ represents the number of nodes that have interacted with node i within the designated time window, and $|Y|$ denotes the total number of previous time windows considered for the calculation. Assuming $Y < I$, the complexity can be simplified to $O(I)$.

Similarly, while computing similarity as shown in Equation 11, the computational complexity can be described as $O(|C| + |I| + |X|) = O(|I|)$. Subsequently, the computational complexity associated with the indirect reputation calculation can be represented as $O(|C| * (|I| + |I|)) = O(|C| * |I|)$.

The computational complexity of the final reputation calculation can be viewed as the sum of the complexities of both direct and indirect reputation calculations. Hence, its computational complexity can be formulated as $O(|I| + |C| * |I|)$, which can be simplified to $O(|C| * |I|)$. Despite resulting in a quadratic complexity, given the small values of both $|I|$ and $|C|$, contemporary and future edge devices can readily handle such computational demands.

2) COMMUNICATION COMPLEXITY

Reputation-related data sharing in the modified zero trust network utilises a Blockchain, as shown in Figure 2. Assuming the size of a message is m bytes, the communication complexity per node in a time window t_y can be approximated as $O(m * |I| * |Y|)$. Despite the quadratic nature of communication complexity, it remains small due to the small values of both $|I|$ and $|Y|$. It is important to note that the communication complexity associated with the Blockchain is not considered here.

Since both computational and communication complexity in reputation calculations are minimal, this suggests that the processing delay and communication latency for reputation computation are low. However, in large-scale deployments where individual nodes interact with a substantial number of other nodes, both computational and communication complexities can increase significantly. This is particularly true as the size of the interaction set $|I|$ and the common set of nodes $|C|$ can be quite large. Such an increase in complexities can directly impact the performance of the trust calculation mechanism.

IV. EXPERIMENT SETUP AND SCENARIOS

To assess the viability of the reputation management system within the proposed modified ZTA, the LEAF simulator [30] was used.

A. CONFIGURATION OF DEVICES

The simulation environment consists of fixed and mobile end devices, edge servers, and the cloud. Table 1 presents the typical configuration parameters of these devices. The computing capability of devices is quantified using Compute Units (CUs), where 1 CU corresponds to 1 million instructions per second. The energy consumption of the cloud is modelled with 0.5 W/CU [30].

TABLE 1. Device configuration.

Device Type	fixed end device	mobile end device	edge server	cloud
Mobile	no	yes	no	no
Speed (m/s)	0	1.5	0	0
Computing capability (CUs)	1	1	400	inf
Maximum power (W)	1.8	1.8	200	-
Static power (W)	0.2	0.2	30	-

B. CONFIGURATION OF COMMUNICATION LINKS

WiFi is used for end devices to connect to the nearest edge server and for edge servers to connect to nearby edge servers. Edge servers use WAN to connect to the cloud. WiFi and WAN bandwidths are set to 1.3×10^9 bps and 100×10^6 bps, respectively [30]. As for the energy models, WiFi employs a value of 300 nJ/bit, while that of WAN is 6000 nJ/bit [31].

C. APPLICATION CONFIGURATION

It is assumed in this simulation that both fixed and mobile end devices are engaged in periodic ML model-inferring tasks. The reputation of the end devices is influenced by the quality of the input data they provide for ML inference, which is evaluated by the edge servers. The quality of the input data is assessed on the basis of the following criteria: *i*) *timeliness*: data is delivered within the specified time frame; and *ii*) *believability*: data is accompanied by a digital signature and the authenticity of the data owner can be verified; The data used in the simulation is sourced from the dataset available at [32] and it incorporates the *disruptive* and *delayed messages* misbehaviour models. In the LEAF simulator, tasks are measured in Compute Units (CUs). In this simulation, both fixed and mobile end devices possess 1 CU each and are tasked with a workload of 100 CUs for ML model training, forcing them to offload their tasks.

D. CONFIGURATION OF MALICIOUS EDGE SERVERS

Our aim is not to propose a new technique for mitigating trust-based attacks. Instead, our focus is on simulating these attacks and assessing the resilience of the proposed trust management system within the modified ZTA. To this extent, the robustness of the trust management system against trust-based security attacks is evaluated by utilising the following configurations for malicious edge servers.

- *MaliciousEdgeType1* (bad-mouthing): Always report negative ratings for benign end devices to lower their reputation intentionally.
- *MaliciousEdgeType2* (ballot stuffing): Always assign high ratings to other malicious end devices to artificially inflate the reputation of a group of malicious nodes.

Note that only malicious behaviours that are logical and relevant in a ZTA are considered in our evaluation. For example, the act of self-promoting by edge servers is deemed irrelevant in a zero trust context as zero trust networks are focused on clients' (i.e., end devices') authentication and authorisation.

E. SIMULATION SCENARIO

The following simulation scenario is used to evaluate the feasibility of the reputation management system for the proposed modified ZTA.

The simulation scenario follows an urban city environment, featuring rectangular blocks that cover an area with 9 crossings, as depicted in Figure 4. Each crossing is equipped with 4 smart traffic lights, that are connected to an edge server. End devices, such as tablets, smartphones, and smart CCTV cameras, establish connections with these edge servers to perform their periodic tasks. Before establishing a connection, an end device takes into account the proximity and reputation of the edge servers. If the available resources are insufficient to execute a particular task, the edge server offloads the task to the cloud for execution. Each edge server keeps a record of its interactions with the end devices in a reputation database. As described in Section IV-C, edge servers mark these interactions as positive or negative according to the quality of data provided by the end devices for ML inference. Edge servers utilise these interactions to calculate the reputation of the end devices, which is then used to determine whether or not to provide services to them.

Table 2 outlines the simulation configuration parameters. The number of fixed end devices and edge servers corresponds to the number of crossings depicted in the city scenario in Figure 4. To illustrate the impact of malicious behaviours, two malicious edge servers were introduced, each exhibiting distinct malicious behaviours (see Section IV-D). The city block size is modelled after Manhattan city [30]. The application period was deliberately set to generate high traffic, facilitating the evaluation under worst-case scenario. The initial reputation value and the reputation threshold were determined based on previous works by [21] and [33]. After the initialization phase, nodes are chosen based on their reputation value, which must exceed the specified threshold (i.e., reputation > 0.5). This approach effectively mitigates whitewashing attacks, as malicious newcomers lack sufficient reputation to interact with other nodes (i.e., their reputation $\ngtr 0.5$). If a node has previously undertaken tasks similar to the current one, its recommendation is weighted at 80% (i.e., $\delta = 0.8$), akin to the value of γ. The remaining parameters, i.e., γ, Z, Y, and time slot length, are adopted from [10].

TABLE 2. Simulation parameters.

Parameter	Value
#devices: fixed, mobile, edge nodes	9, 15, 9
#malicious edge nodes: type1, type2	1 each
simulation block size (W x H)	274 m x 80 m
application period	5 seconds
initial reputation	0.5
time slot length	60 seconds
γ, δ, Z, Y	0.5, 0.8, 0.8, 10
reputation threshold	0.5
initialisation time	5 minutes
simulation time	4 hours

V. PERFORMANCE EVALUATION

In this section, the simulation results are evaluated to determine the feasibility of the proposed modified ZTA, as illustrated in Figure 2.

To assess the stability of the simulations, the confidence interval (CI) of the simulation runs is first evaluated. After conducting 12 iterations, the 95% CI for the application delay in the two scenarios, without and with reputation management, are 2.960 ± 0.069 seconds and 2.964 ± 0.102 seconds, respectively. Since the CIs are small compared to the mean in both cases, it can be inferred that the impact of random seeds on the variability of the results is negligible.

Examining the influence of the reputation management mechanism on application performance is crucial. Figure 5 compares the average application delay for fixed and mobile end devices with and without reputation management. As depicted in the figure, the hourly disparity in application delay between the scenarios without and with reputation management is minimal. In addition, the overall impact on

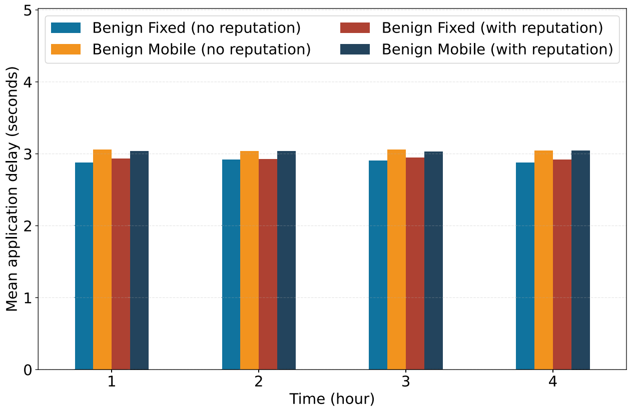

FIGURE 5. Impact of reputation management on application delay.

application delay can be measured as a minor increase of 37 milliseconds for fixed end devices and a slight decrease of 13 milliseconds for mobile end devices, indicating that reputation management does not noticeably affect application performance. In particular, mobile devices experience reduced application delay when using the reputation management system, primarily because they are exposed to edge servers with higher reputations than fixed devices.

Within the zero trust framework, edge servers maintain records of end devices' reputations, enabling the policy engine to determine access rights. During the simulation, edge servers assess the quality of data provided by the end devices for ML inference to compute their reputation. However, it is essential to be aware that malicious nodes can potentially disrupt the reputation management process. Figure 6 presents the percentages of successful execution of tasks by edge servers. In the absence of a reputation management scheme, distinguishing between the two types of end devices becomes difficult, which is highly undesirable in the context of zero trust. However, the introduction of the reputation management system reveals a distinct ability to identify malicious end devices from benign ones, with the average successful task executions during the simulation period being 19.8 % for malicious devices and 97.5 % for benign ones. This effectively mitigates the threats posed by malicious end devices.

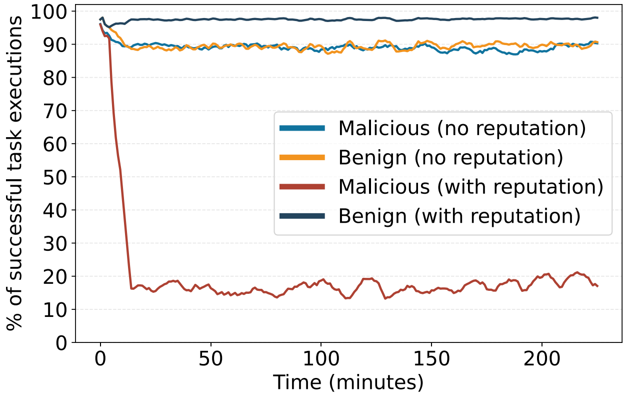

FIGURE 6. Successful task completions by malicious and benign edge servers.

Figure 7 depicts the performance of reputation management in the presence of malicious edge servers (bad-mouthing and ballot stuffing) and malicious end devices (non-timely and untrustworthy data). The edge servers calculate the reputation towards end devices, and the figure illustrates the average reputation of similar edge servers, that is, bad-mouthing, ballot stuffing, and benign, towards similar end devices, i.e., D1, D2, and benign. Notably, the simulation includes one device for each malicious behaviour. (see Table 2). It is important to note that there are missing data points for the two malicious edge servers. This is because the corresponding malicious end devices did not interact with them. The impact of negative ratings from the bad-mouthing edge server towards fixed and mobile benign end devices is evident, as their reputations are lower compared to the D1 malicious end devices. Additionally, none of the end devices received a reputation that exceeded the specified reputation threshold of 0.5. Conversely, the ballot stuffing edge server deliberately inflates the reputation of end devices, as expected. As a non-malicious edge server, the reputation management scheme successfully distinguishes between malicious and benign end devices. This distinction is evident in the figure, where the non-malicious edge servers assign lower reputations to malicious end devices compared to benign ones.

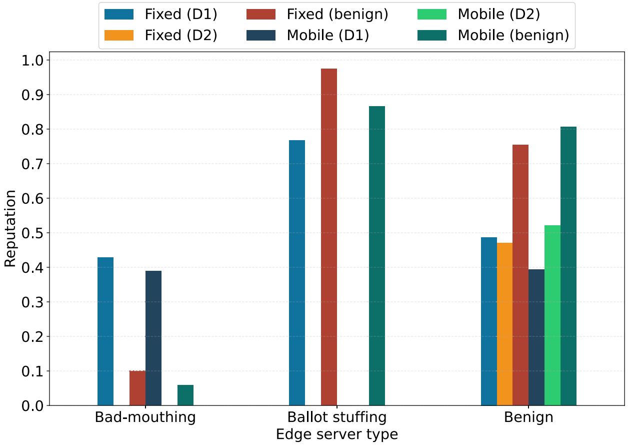

FIGURE 7. Reputation of edge servers towards end devices. Here, D1 represents the non-timely sharing of data by the end devices, while D2 represents the sharing of untrustworthy data. Both D1 and D2 are considered to be malicious activities (see Section IV-C).

The energy cost per user authentication and authorisation request in a zero trust network is influenced by various factors, such as the authentication method and the hardware/infrastructure utilised. Due to the complexity and various implementations of zero trust, providing an exact energy cost per request is a challenge. However, for this analysis, the following assumptions are made: *i*) X.509 certificates are generated using the ECDSA-256 cryptographic algorithm [34], the average energy consumption for key generation, signature generation, and verification is estimated to be 2.34 J per request [35]. *ii*) All communication is encrypted with AES-128 and the associated energy cost per communication is 0.67 J [35]. *iii*) JSON Web Tokens (JWTs) are used to grant short-term access to clients, and Base64-URL encoding/decoding is used. The energy cost associated with this process is negligible. *iv*) one access request involves three communications: access request from the server, service reconfiguration to allow the client access, and access grant through a leased access token (i.e., a JWT). Consequently, the total energy cost per access request is 4.35 J. *v*) To estimate the energy consumption associated with the reputation calculations (i.e., arithmetic operations), the work in [36] was followed. Note that authors are aware that employing X.509 certificates in scenarios involving multiple domains and stakeholders is not practical. The management of these certificates, including certificate issuance, revocation, and renewal, poses significant challenges in maintaining security and trust. Moreover, establishing a trusted PKI infrastructure may require interoperability and trust agreements among various parties. Despite these challenges, for the purpose of evaluating the worst-case scenario regarding energy consumption, it is assumed that the nodes use X.509 certificates for authentication.

Figure 8 illustrates the performance of reputation management in terms of energy savings. Mobile end devices consume more energy than fixed ones due to their higher numbers. With the reputation management scheme, mobile end devices saved 203.9 kWh over 4 hours, while fixed end devices saved 12.7 kWh. This is because there are fewer fixed devices (see Table 2), and they have fewer options for selecting an edge server. However, mobile end devices achieve higher per-node energy savings.

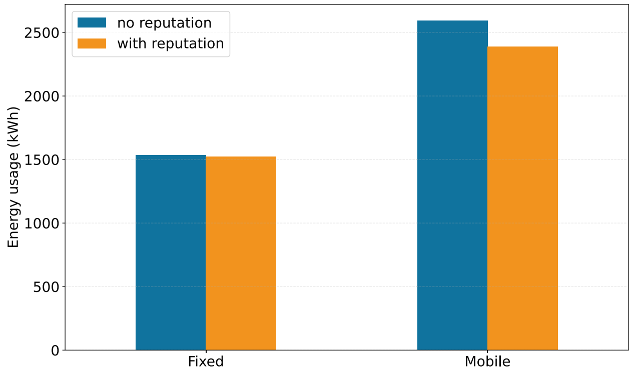

FIGURE 8. **The average energy consumption with and without reputation management.**

Finally, it should be noted that the observations presented in this section do not reject our hypothesis stated in Section I.

VI. CONCLUSION AND FUTURE RESEARCH DIRECTIONS

This paper presented a novel zero trust framework for the next generation of IoT services and explained its key components. In particular, the performance-based trust assessment was discussed in depth to determine its feasibility with the proposed modified zero trust framework. The paper assessed the computational and communication complexity of the performance-based trust assessment and further evaluated its performance in different simulation scenarios. The results demonstrated that the performance-based trust assessment mechanism is feasible for resource-constrained edge computing environments. It is important to highlight the potential scalability challenges of our technique in large-scale deployments. As discussed in Section III-E1, the computational complexity of the proposed performance-based trust evaluation mechanism grows with the number of devices. Thus, it is reasonable to anticipate performance degradation in extensive deployments where individual nodes interact with a substantial number of others.

This study did not evaluate the performance of the reputation management mechanism for newcomers. Therefore, future works incorporate a mechanism for newcomers involving a dummy task, the outcome of which determines their fate (similar to the approach in [33]). Subsequently, the reputation management mechanism will be evaluated with newcomers, particularly malicious ones, to gauge the robustness of our approach against whitewashing attacks. Further future work includes investigating contextual and behavioural pattern-based anomaly detection using ML models to enhance the security of the zero-trust network. Additionally, exploring the integration of performance-based trust assessment and ML-based anomaly detection with Blockchain for data sharing among distributed nodes is a potential avenue for further research. This paper explored the feasibility of a trust assessment mechanism designed for our conceptualised ZTA (see Figure 2). An experimental comparison between our proposed conceptualised ZTA, the NIST ZTA, and other ZTAs is on our agenda for future research, following the implementation and testing of the individual components such as the updated reputation management mechanism with newcomers, contextual and behavioural pattern-based anomaly detection using ML models, and integration of Blockchain in the proposed ZTA. Furthermore, a testbed is being designed and developed that encompasses the entire device-edge-cloud continuum to substantiate the efficacy of our modified ZTA.

REFERENCES

[1] Y. He, D. Huang, L. Chen, Y. Ni, and X. Ma, "A survey on zero trust architecture: Challenges and future trends," *Wireless Commun. Mobile Comput.*, vol. 2022, pp. 1–13, Jun. 2022.

[2] D. D'Silva and D. D. Ambawade, "Building a zero trust architecture using kubernetes," in *Proc. 6th Int. Conf. for Converg. Technol. (I2CT)*, Apr. 2021, pp. 1–8.

[3] N. Ghate, S. Mitani, T. Singh, and H. Ueda, "Advanced zero trust architecture for automating fine-grained access control with generalized attribute relation extraction," in *Proc. Int. Conf. Emerg. Technol. Commun.*, vol. 68, 2021, pp. 1–28.

[4] K. Ramezanpour and J. Jagannath, "Intelligent zero trust architecture for 5G/6G networks: Principles, challenges, and the role of machine learning in the context of O-RAN," *Comput. Netw.*, vol. 217, Nov. 2022, Art. no. 109358.

[5] J. Kindervag, "Build security into your network's DNA: The zero trust network architecture," Forrester Res., 2012.

[6] C. DeCusatis, P. Liengtiraphan, A. Sager, and M. Pinelli, "Implementing zero trust cloud networks with transport access control and first packet authentication," in *Proc. IEEE Int. Conf. Smart Cloud (SmartCloud)*, Nov. 2016, pp. 5–10.

[7] J. Olsson, A. Shorov, L. Abdelrazek, and J. Whitefield, "5G zero trust—A zero-trust architecture for telecom," *Ericsson Technol. Rev.*, vol. 2021, no. 5, pp. 2–11, May 2021.

[8] P. Phiayura and S. Teerakanok, "A comprehensive framework for migrating to zero trust architecture," *IEEE Access*, vol. 11, pp. 19487–19511, 2023.

[9] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero trust architecture," Nat. Inst. Sci. Technol. (NIST), Tech. Rep. NIST Special Publication 800-207, 2020.

[10] J. Kang, Z. Xiong, D. Niyato, S. Xie, and J. Zhang, "Incentive mechanism for reliable federated learning: A joint optimization approach to combining reputation and contract theory," *IEEE Internet Things J.*, vol. 6, no. 6, pp. 10700–10714, Dec. 2019.

[11] J. Yuan and X. Li, "A reliable and lightweight trust computing mechanism for IoT edge devices based on multi-source feedback information fusion," *IEEE Access*, vol. 6, pp. 23626–23638, 2018.

[12] X. Huang, R. Yu, J. Kang, and Y. Zhang, "Distributed reputation management for secure and efficient vehicular edge computing and networks," *IEEE Access*, vol. 5, pp. 25408–25420, 2017.

[13] J. Liang, M. Zhang, and V. C. M. Leung, "A reliable trust computing mechanism based on multisource feedback and fog computing in social sensor cloud," *IEEE Internet Things J.*, vol. 7, no. 6, pp. 5481–5490, Jun. 2020.

[14] M. Al-khafajiy, T. Baker, M. Asim, Z. Guo, R. Ranjan, A. Longo, D. Puthal, and M. Taylor, "COMITMENT: A fog computing trust management approach," *J. Parallel Distrib. Comput.*, vol. 137, pp. 1–16, Mar. 2020.

[15] R. Latif, M. U. Ahmed, S. Tahir, S. Latif, W. Iqbal, and A. Ahmad, "A novel trust management model for edge computing," *Complex Intell. Syst.*, vol. 8, no. 5, pp. 3747–3763, Oct. 2022.

[16] Y. Liu, X. Hao, W. Ren, R. Xiong, T. Zhu, K. R. Choo, and G. Min, "A blockchain-based decentralized, fair and authenticated information sharing scheme in zero trust Internet-of-Things," *IEEE Trans. Comput.*, vol. 72, no. 2, pp. 501–512, Feb. 2023.

[17] Y. Feng, Z. Zhong, X. Sun, L. Wang, Y. Lu, and Y. Zhu, "Blockchain enabled zero trust based authentication scheme for railway communication networks," *J. Cloud Comput.*, vol. 12, no. 1, pp. 1–24, Apr. 2023.

[18] X. Deng, J. Liu, L. Wang, and Z. Zhao, "A trust evaluation system based on reputation data in mobile edge computing network," *Peer Peer Netw. Appl.*, vol. 13, no. 5, pp. 1744–1755, Sep. 2020.

[19] G. Fortino, L. Fotia, F. Messina, D. Rosaci, and G. M. L. Sarné, "Trusted object framework (TOF): A clustering reputation-based approach using edge computing for sharing resources among IoT smart objects," *Comput. Electr. Eng.*, vol. 96, Dec. 2021, Art. no. 107568.

[20] L. Rui, S. Yang, Z. Gao, W. Li, X. Qiu, and L. Meng, "Smart network maintenance in edge cloud computing environment: An allocation mechanism based on comprehensive reputation and regional prediction model," *J. Netw. Comput. Appl.*, vol. 198, Feb. 2022, Art. no. 103298.

[21] L. Guo, H. Yang, K. Luan, Y. Luo, L. Sun, and X. Zheng, "A trust management model based on mutual trust and a reward-with-punishment mechanism for cloud environments," *Concurrency Comput., Pract. Exper.*, vol. 33, no. 16, pp. 1–18, Aug. 2021.

[22] W. Li, J. Cao, K. Hu, J. Xu, and R. Buyya, "A trust-based agent learning model for service composition in mobile cloud computing environments," *IEEE Access*, vol. 7, pp. 34207–34226, 2019.

[23] G. Fortino, L. Fotia, F. Messina, D. Rosaci, and G. M. L. Sarnè, "A social edge-based IoT framework using reputation-based clustering for enhancing competitiveness," *IEEE Trans. Computat. Social Syst.*, vol. 10, no. 4, pp. 1–10, Sep. 2022.

[24] R. N'goran, J.-L. Tetchueng, G. Pandry, Y. Kermarrec, and O. Asseu, "Trust assessment model based on a zero trust strategy in a community cloud environment," *Engineering*, vol. 14, no. 11, pp. 479–496, 2022.

[25] M. Hao, D. Ye, S. Wang, B. Tan, and R. Yu, "URLLC resource slicing and scheduling for trustworthy 6G vehicular services: A federated reinforcement learning approach," *Phys. Commun.*, vol. 49, Dec. 2021, Art. no. 101470.

[26] N. Kandhoul, S. K. Dhurandher, and I. Woungang, "T_CAFE: A trust based security approach for opportunistic IoT," *IET Commun.*, vol. 13, no. 20, pp. 3463–3471, Dec. 2019.

[27] K. Kalkan and K. Rasmussen, "TruSD: Trust framework for service discovery among IoT devices," *Comput. Netw.*, vol. 178, Sep. 2020, Art. no. 107318.

[28] S. Dhelim, N. Aung, M. T. Kechadi, H. Ning, L. Chen, and A. Lakas, "Trust2 Vec: Large-scale IoT trust management system based on signed network embeddings," *IEEE Internet Things J.*, vol. 10, no. 1, pp. 553–562, Jan. 2023.

[29] A. Jøsang, "A logic for uncertain probabilities," *Int. J. Uncertainty, Fuzziness Knowl.-Based Syst.*, vol. 9, no. 3, pp. 279–311, Jun. 2001.

[30] P. Wiesner and L. Thamsen, "LEAF: Simulating large energy-aware fog computing environments," in *Proc. IEEE 5th Int. Conf. Fog Edge Comput. (ICFEC)*, May 2021, pp. 29–36.

[31] S. Kakati and R. Deka, "Computational and adaptive offloading in edge/fog based IoT environments," in *Proc. 2nd Int. Conf. Intell. Technol. (CONIT)*, Jun. 2022, pp. 1–6.

[32] J. Kamel, M. Wolf, R. W. van der Hei, A. Kaiser, P. Urien, and F. Kargl, "VeReMi extension: A dataset for comparable evaluation of misbehavior detection in VANETs," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2020, pp. 1–6.

[33] S. Bharti and A. McGibney, "CoRoL: A reliable framework for computation offloading in collaborative robots," *IEEE Internet Things J.*, vol. 9, no. 19, pp. 18195–18207, Oct. 2022.

[34] FIDO Alliance. (2023). *User Authentication Specifications Overview*. [Online]. Available: https://fidoalliance.org/specifications/

[35] S. Seys and B. Preneel, "Power consumption evaluation of efficient digital signature schemes for low power devices," in *Proc. IEEE WiMob Int. Conf. Wireless Mobile Comput., Netw. Commun.*, Jul. 2005, pp. 1–29.

[36] M. Horowitz, "1.1 Computing's energy problem (and what we can do about it)," in *Proc. IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers (ISSCC)*, Feb. 2014, pp. 10–14.

INDIKA DHANAPALA (Member, IEEE) received the Ph.D. degree in electronic engineering from Munster Technological University, Cork, Ireland, in 2022. Since 2022, he has been a Postdoctoral Research Fellow with the Nimbus Research Centre, Munster Technological University. His research interests include trusted distributed systems, particularly in the edge cloud continuum and in open-RAN architecture, and the Internet of Things.

SOURABH BHARTI (Member, IEEE) received the Ph.D. degree in information technology from Indian Institute of Information Technology and Management, Gwalior, India, in 2018. Currently, he is with the Nimbus Research Centre, Munster Technological University, Cork, Ireland. He is also an Associate Investigator within the nationally funded CONNECT Centre, Ireland. His research interests include applied artificial intelligence, machine learning, and the Internet of Things.

ALAN MCGIBNEY (Member, IEEE) received the Ph.D. degree in electronic engineering from Cork Institute of Technology, Ireland, in 2008. He is currently a Group Lead of IoT systems and user interaction with the Nimbus Centre, Munster Technological University, with a particular focus on the areas of distributed data management, tools and service development, software architectures, and trusted digital platforms. He leads and coordinates a number of EU funded and national projects in this area and is an SFI funded investigator as part of the CONNECT Centre for future networks contributing to industrial IoT and wireless communications research activities. His research interests include the Internet of Things (IoT) and cyber physical systems.

SUSAN REA (Senior Member, IEEE) is currently the Centre Director with the Nimbus Centre, Munster Technological University. She is a Principal Investigator with the SFI CONNECT Centre for Future Networks. Her current research interests include the Internet of Things and cyber physical systems, specifically embedded infrastructure management using distributed ledger technology and cybersecurity for trusted large scale next generation networks. She is a member of the INATBA's Academic Advisory Body and the Blockchain Ireland Skills, Education and Innovation Working Group and being a member (and former the Vice Chair) of the IEEE U.K. and Ireland Blockchain Group.