A fuzzy reputation system for Radio Access Network sharing

Pierre-Marie Lechevalier, Romaric Ludinard, Yann Busnel, Géraldine Texier

A fuzzy reputation system for Radio Access Network sharing

Pierre-Marie Lechevalier
IMT Atlantique / IRISA
Rennes, FRANCE
pierre-marie.lechevalier@imt-atlantique.fr

Yann Busnel
IMT Nord Europe / IRISA
Lille, FRANCE
yann.busnel@imt-nord-europe.fr

Romaric Ludinard
IMT Atlantique / IRISA
Rennes, FRANCE
romaric.ludinard@imt-atlantique.fr

Géraldine Texier
IMT Atlantique / IRISA
Rennes, FRANCE
geraldine.texier@imt-atlantique.fr

***Abstract*—5G network slicing allows the coexistence of multiple virtualized networks on the same infrastructure. Leveraging this concept, it becomes possible to create marketplaces where Infrastructure Providers (InPs) lease network resources to different Mobile Virtual Network Operators (MVNOs) while accommodating their specific Quality of Service (QoS) requirements. In addition to the cost criteria, an MVNO choosing between several InP might be interested in evaluating the InPs actual capacity to deliver the expected Service Level Agreement (SLA). Existing 5G literature on trust evaluation is often based on blockchain. While this technology offers transparency and auditability, the inherent consensus mechanism often brings additional costs.**

In this paper, we propose a distributed reputation system based on fuzzy logic that can provide a robust and dynamic estimation of an InP behavior while respecting the subjective requirements of MVNOs. We evaluate this system in a Radio Access Network (RAN) sharing simulation and we show that it can redirect MVNO to InP that are capable of meeting their needs. We further test different trust decay strategy in order to find one that is able to both quickly react to a network outage and forgive InP once the punctual outage is solved.

***Index Terms*—reputation system, 5G, ran-sharing, trust, similarity**

I. Introduction

Fifth generation radio networks needs to address the growing demand for user data exchange and the arrival of new vertical use cases with their specific requirements. Meanwhile, it also aims at simplifying and automating network management. Network virtualization appears as a key enabler to reach this goal. Indeed, in 5G, Network Function (NF) follow strict specifications and are virtualized to run on customer of the shelf equipment. NF orchestration is standardized as well in a framework called Management and Orchestration (MANO). This virtualization enables the concept of slicing, in which multiple virtual infrastructures coexist in isolation while sharing a common physical infrastructure. The benefits of this approach are twofold. Firstly, slices can be fine-tuned to better match the specific requirements of each vertical use case; secondly, the better interoperability opens the way for more flexible infrastructure sharing between multiple tenants [1].

There are many reasons for sharing network infrastructure: reducing capital and operational expenditures, complying with the law, or improving coverage in areas where deploying several dedicated networks isn't economically sound [2]. These benefits are especially apparent in the Radio Access Network (RAN) which concentrate most of the cost and energy consumption in the network. As a result, RAN sharing recently received a great deal of attention [3]–[8]. Also, despite being less dynamic than what is currently envisioned, some RAN sharing scenario between operators are already in place. For instance, some telecom operators already share their physical infrastructures in the least economically profitable areas. This kind of agreement where only physical infrastructure is shared is called Multi-Operator Core Network (MOCN). In opposition, the case where frequency bands are additionally shared is called Multi-Operator Radio Access Network (MORAN) [2]. In spite of its appealing features, network infrastructure sharing requires Mobile Virtual Network Operators (MVNOs), *i.e*. tenants renting resources, to trust the Infrastructure Providers (InPs) they are leasing from [5]. Indeed, a breach of Service Level Agreement (SLA) from the InP will directly affect the MVNO clients [9]. Also, it can be in the InP best interest to over-commit his infrastructure to multiple MVNO in order to maximize his profits [6]. A trust mechanism to monitor the good behavior of InPs is thus necessary. Ideally, this mechanism should be decentralized in order to avoid reliance on a third party. Reputation systems can dynamically capture and distribute feedback from past interaction in order to guide future interactions choice [10]. While only few of them exist in the 5G Network Service (NS) sharing use case [11], [12], reputation mechanisms are regarded as a possible candidate to manage the trust issue in 5G resource sharing [12]. In this paper, we design a distributed fuzzy logic based reputation system that is able to:

- handle the subjectivity of the SLA requirements from MVNO that have different use cases;

This work is part of the Beyond5G project, which is funded by Banque Publique d'Investissement (BPI) as part of the *France Relance* investment plan. The authors are thankful to Hélène Le Bouder for her advising and proofreading.

- provide a dynamic reputation score that can match an InP outage;
- penalize InPs that present an oscillatory behavior in order to maximize their profit.

We extensively test the proposed reputation system in a shared RAN simulation against varying abusive InP behaviors.

The rest of this paper is organized as follows. Section II presents some necessary background in 5G use case and resource sharing. Additionally, it presents related works on trust and reputation and their limitations. Section III presents our reputation system design and the underlying motivations for our choices. We then present the evaluation use case and implementation choices in Section IV and provide an evaluation of our approach in Section V. Finally, we conclude and lay out future works in Section VI.

II. Background and related work

This section provides some background on the main 5G use cases, RAN resource sharing, and the existing works regarding trust and reputation in 5G resource sharing.

A. 5G use cases

One of the benefit of network virtualization is to enable a potentially large number of use cases to coexist while having radically different requirements. For the sake of clarity and without losing generality, we will focus in this paper on only three of them:

enhanced Mobile BroadBand (eMBB): offers a better mobile data user experience by increasing bandwidth and reducing latency,

Ultra-Reliable Low-Latency Communications (URLLC): targets critical services and provides an increased reliability (*i.e.* less packet drop) and very low latency,

massive Machine Type Communication (mMTC): targets industry and Internet of Things (IoT) devices with an aim to offer an energy efficient connection to numerous devices.

The variety of use cases makes resource sharing more complex, one InP might provide NS relevant for eMBB but inadequate for URLLC. Furthermore, as MVNOs have potentially different requirements, transferring the opinion that an MVNO have on a given InP to another MVNO requires some attention.

B. RAN resource sharing

Fifth generation advances in virtualization and interoperability has led to a growing interest in RAN resource sharing [3]–[5], [7], [8]. For the rest of this paper, and inline with [4], [5], the entity furnishing a NS will be referred to as InP and the entity benefiting from the service will be referred to as MVNO. Please note that while we make this distinction for reading clarity, in practice an MVNO could own some network resources and could even be an InP granted that he subleases some network resources to other MVNOs. Different levels of network resources can be shared. On the one hand, Ou *et al.* [4] envision a multi-tenant RAN that follow the Open-RAN (O-RAN) functional split where each MVNO own their Radio Units (RUs) but rent Distributed Units (DUs) and Centralized Units (CUs) from an InP. On the other hand, Vilà *et al.* [8] propose a capacity sharing scheme where the InP directly supplies Resource Blocks (RBs) to the different tenants. In this last scheme, tenants are isolated through the instantiation of multiple RAN slice instance (RSI) on the InP infrastructure. Both of these works define resource allocation strategies so that InPs can adequately address MVNOs SLA requirements while maximizing the number of accepted requests and thus their profit. Pérez-Romero *et al.* [7] also propose a resource allocation strategy, they even go one step further in profit maximization as they consider the possibility of breaching MVNOs SLA requirements. In a profit driven context where InPs are incited to over-commit the same physical resource to multiple MVNOs [6], choosing trustworthy InPs is a reasonable consideration for MVNOs that wishes to guarantee their end-users SLA.

C. Trustworthy resource sharing

Benefits of SLA enforcement mechanism in a resource sharing scenario have been extensively studied [6], [9], [12], [13]. These mechanisms rely either on a blockchain [6], [9], [13] or on a reputation system [12].

1) SLA enforcement using blockchain: Many of these contributions propose solutions based on blockchain and smart contracts [6], [9], [12], [13]. Faisal *et al.* [9] and Boateng *et al.* [6] for instance propose to log past transactions onto a permissioned blockchain and to decide reward and punishment using a smart-contract. The benefit of this approach is to provide an immutable, transparent and auditable list of existing transactions [13]. However, it comes at a cost in terms of convergence time and scalability [13]. Also, while blockchain technologies offer an immutable record, there is no guarantee that the data is trustworthy as it is committed by the InP. While it's possible to circumvent this issue by placing the monitoring functions in a Trusted Execution Enclave (TEE) [9], these enclaves suffer from their own security issues and are not present on all network equipments.

2) Reputation system as an alternative: In this paper, we defend an alternative solution where SLA informations are gathered from MVNO's and combined using a reputation system. Reputation systems subjectively assess participants' ability to perform a task based on the feedback from past interactions [10]. In the 5G resource sharing context, interactions are resource leased from an InP to an MVNO and the feedback is the appreciation of the outcome of this transaction by the MVNO. A failed transaction could be due to a complete outage or just a failure from the InP to comply with the predefined SLA. Their process can be distributed among different participants to remove dependence on trusted third-party [14].

Valero *et al.* [12] propose a reputation system based on SLAs compliance for a 5G service marketplace. One of the added value of this work is to propose a continuous update module that can continue to track trust on ongoing transactions. This continuous monitoring allows to finely decide when the provided Quality of Service (QoS) levels are not meeting the

initial SLA anymore and the transaction should stop. Their trust model is an adaptation from PeerTrust [15], it aggregates the subjective satisfaction of several participants to obtain the final reputation score. Although this is a common approach, we believe that aggregating the satisfaction of MVNOs with potentially very different use case and requirements is not appropriate (see Section II-A). In the approach we are proposing, the different MVNOs share objective measurements on the provided QoS and each MVNO can then construct his own opinion depending on its subjective SLA requirements.

D. Attacks on reputation

The feedback aggregated by the reputation system comes from multiple users who are not necessarily trusted. Moreover, InPs knowing that they are being monitored by a reputation system might adapt their behavior to bias the system in their favor. It is thus necessary to protect the reputation system against attacks that target it. Koutrouli *et al.* [14] suggest a taxonomy that regroup these attacks in three main categories for peer-to-peer networks. Let's review these categories and see how they apply to the 5G context.

Unfair recommendations: participants in the reputation system can spread false feedback to unfairly increase or decrease the reputation score of other participants. This problem stays relevant for 5G applications, common mitigation such as feedback normalization and filtering are nonetheless applicable and should be included in the reputation system design.

Identity management attacks: it regroups (a) *sybils—i.e.* multiple participants created and controlled by a single entity; (b) *whitewashing—i.e.* discarding a low reputation through an identity reset, and (c) *impersonation*— wherein a participant is able to forge feedback in place of other participants. While *identity management* can be hard to enforce when combined with privacy [16], 5G participants are moral entities and their interactions don't need to be private. It is thus possible to address *impersonation* by authenticating the participants and signing the feedback requests. Additionally, access to the system can be controlled with an entry price to dissuade *whitewashing* and *sybils* attacks.

Inconsistent behaviors: participants might dynamically adapt their behaviors to exploit the reputation system and maximize their profit. An InP could for instance initially provide a service of good quality to build up reputation and then suddenly or punctually over-commit his infrastructure to make the most out of his reputation. Another possibility would be to provide services that oscillate in quality to limit the infrastructure overhead while maintaining a good reputation overall. This last case is named *oscillatory behavior* [14]. These attacks are commonly addressed through the combination of historical contribution in a decay function [14]–[16].

TABLE I: Notations

Notation	Description
$\mathbb{M}$	Set of all Mobile Virtual Network Operators
m_i	Mobile Virtual Network Operator i
$\mathbb{P}$	Set of all Infrastructure Providers
p_k	Infrastructure Provider k
$\overrightarrow{f_{i,k}}$	Vector with all feedbacks emitted by m_i on p_k
$f^n_{i,k}$	n$^{\text{th}}$ feedback emitted by m_i on p_k
$\lvert\overrightarrow{f_{i,k}}\rvert$	Number of transactions present in the feedback vector by m_i on p_k
$F_{i,*}$	Matrix with $\overrightarrow{f_{i,k}}\ \forall k \in \mathbb{P}$
$\sigma_{i,j}$	Similarity of feedbacks between m_i and m_j
$Cr_{i,j}$	Credibility of m_j feedbacks as perceived by m_i
$T_{i,k}$	Trust of p_k as computed by m_i based on its own past transaction.
$R_{i,k}$	Reputation of p_k as computed by m_i considering its own past transaction and external feedback.

III. Architecture

This section details the reputation system architecture, the interactions, and inputs of the different reputation bricks. We first present in Section III-A the decay functions used to address *inconsistent behaviors* before outlining in Section III-B the proposed similarity mechanism that aims to limit the effect of *unfair recommendations*. We finally detail our choice of a fuzzy reputation system and explain how fuzzy logic works in Section III-C.

A. Decaying reputation

The purpose of a decay function in a reputation system is usually twofold. First, it aims to keep trust information up to date with the latest developments and, secondly, it keeps tracks of the past transactions in order to penalize *inconsistent behaviors*. Moreover, the RAN sharing context brings two additional constraints.

Constraint 1: in case of a network failure from an InP, the reputation should adjust as quickly as possible to warn the MVNOs likely to choose it.

Constraint 2: as opposed to peer-to-peer networks, where many candidates for interaction are available, there are potentially only a few InPs. It is thus not acceptable to permanently affect the reputation of an InP for a single involuntary punctual failure.

In this paper, we compare two common decay function from the literature: *exponential decay* [17] and *adaptive time window* [15].

1) Exponential decay: *exponential decay* gradually decreases the importance of older transactions to make newer transactions relatively more important. The rate at which older transactions are gradually forgotten depends on a parameter $\lambda \in [0, 1]$. A lower λ give more importance to recent transactions, the minimum being $\lambda = 0$ where only the last transaction is taken into account. Conversely, a higher λ even the transactions weight until $\lambda = 1.0$ where all transactions are equal. Equation (1) exposes how the feedback from past transactions is aggregated into a single trust value $T_{i,k}$.

$$T_{i,k} = \sum_{n=1}^{|\overrightarrow{f_{i,k}}|} f_{i,k}^n \cdot \lambda^{|\overrightarrow{f_{i,k}}|-n} \quad (1)$$

In this approach, the λ parameter arbitrates between two distinct goals: the duration for which past opinions remain relevant and the speed at which behavioral change will be taken into account. **Constraint 1** leads to choose small λ, opening the door to *inconsistent behaviors* such as *oscillatory* abuse whose detection require longer records.

2) Adaptive time window: Xiong *et al.* [15] propose a decay strategy that aims to address this issue by leveraging two sliding windows with different sizes. A first trust score $Ts_{i,k}$ is computed on the window $\overrightarrow{fs_{i,k}}$, a small subset of $\overrightarrow{f_{i,k}}$ containing only the more recent transactions. Since this window only considers the last few transactions, $Ts_{i,k}$ can closely match changes in behavior. A second trust score $Tl_{i,k}$ is computed on the window $\overrightarrow{fl_{i,k}}$, a larger subset of $\overrightarrow{f_{i,k}}$. The longer memory period of $Tl_{i,k}$ makes it less prone to oscillatory abuse. As seen in eq. (2), by default $T_{i,k} = Tl_{i,k}$, $T_{i,k}$ only becomes equal to $Ts_{i,k}$ when the small window is worst by a tunable ε margin.

$$T_{i,k} = \begin{cases} Ts_{i,k} & \text{if } Tl_{i,k} - Ts_{i,k} > \varepsilon \\ Tl_{i,k} & \text{otherwise} \end{cases} \quad (2)$$

B. Similarity

Reputation is built by incorporating indirect trust opinions to local trust opinion computed on previous direct interactions. Indeed, by leveraging the opinion from his peers, an MVNO m_i can access more complete and potentially more recent knowledge than by relying solely on his own past interactions. It's also a way to get information on an InP with whom m_i had no interaction.

However, to leverage indirect opinions, the system needs to limit the effect of *unfair recommendations* that could artificially raise or downgrade the ratings of an InP. Indeed, we introduce, for sake of clarity, a semantic distinction between InPs and MVNOs. However, in the general case, a MVNO can also be an InP. When it occurs, MVNOs are thus judge and party, thus creating incentives to provide unfair recommendations on competitors.

When m_i takes into account $\overrightarrow{f_{j,k}}$, *i.e.* m_j feedback on p_k, m_i first computes $\sigma_{i,j}$ referring to the similarity between himself and m_j. $\sigma_{i,j}$ is the similarity between m_i and m_j. Similarly to [15], we compute it based on the proximity of $F_{i,*}$ and $F_{j,*}$ which are m_i and m_j direct observations on all the InPs. Let $I(i)$ denote the set of InP that m_i interacted with and let $I(i,j)$ be $I(i) \cap I(j)$, $\sigma_{i,j}$ is detailed in eq. (3).

$$\sigma_{i,j} = 1 - \sqrt{\frac{\sum_{k \in I(i,j)} \left(\frac{1}{|\overrightarrow{f_{i,k}}|} \sum_{n=1}^{|\overrightarrow{f_{i,k}}|} f_{i,k}^n - \frac{1}{|\overrightarrow{f_{j,k}}|} \sum_{n=1}^{|\overrightarrow{f_{j,k}}|} f_{j,k}^n \right)^2}{|I(i,j)|}} \quad (3)$$

The underlying intuition of $\sigma_{i,j}$ is that if the opinions of two participants on their common InPs is close, they can be confident in the reliability of the feedback they exchange. Equation (4) depict how $\sigma_{i,j}$ can be used by m_i to compute the credibility $Cr_{i,j}$ of m_j relative to other participants. $Cr_{i,j}$ can then be used to aggregate the received indirect feedback together.

$$Cr_{i,j} = \frac{\sigma_{i,j}}{\sum_{l \in \mathbb{M}} \sigma_{i,l}} \quad (4)$$

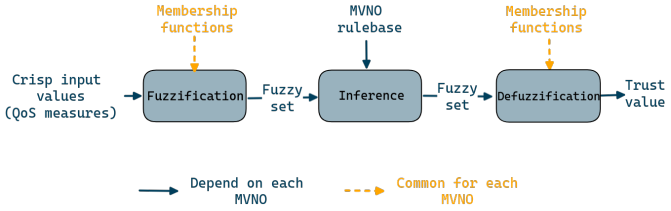

Fig. 1: *Fuzzy logic process with an emphasis on what is shared and what is specific to each MVNO.*

C. Fuzzy aggregation

There are several approaches to trust modelling [15], [18], [19]. Among them, fuzzy logic offers a high degree of granularity and flexibility in decision-making [19]. In particular, it is possible to draw conclusion on imprecise data, or, in our case, on data that can have some legitimate variation from one observation to another. Furthermore, it is possible to combine several input variables, which is relevant since QoS is assessed with several metrics. The different steps used by the fuzzy logic process to compute a score from QoS measurements are depicted in Figure 1. Fuzzy logic associates a set of observations, called *crisp values*, to *terms* that qualify them using *membership functions*. In the *membership function* exposed in Figure 2, the *crisp value* is a QoS measurement, in this case a latency of 4ms. The *terms* are the different appreciation labels *excellent*, *neutral*, *bad*, ... A *membership function* associates a crisp value to the adequate *terms* forming a *fuzzy set*. In this case 4ms is not completely excellent or good but somewhere in between. In this paper, in addition to latency, we use bandwidth and packet loss as crisp values, each having their own membership functions.

The *fuzzy set* obtained can then be used to infer the quality of the InP. This is done using a *rule base* that will combine the *terms* from the input fuzzy set, called antecedents, to terms from an output fuzzy set, called consequent. The *rule base* combines antecedents together using logical sign, *e.g.* AND or OR into a consequent. One of the rule for a MVNO providing *URLLC* services to its client could be: *if (latency is excellent OR latency is good) AND packet drop is excellent, THEN the InP is very trustworthy*.

While the fuzzification membership functions should be the same for every participant, it is possible for each participant to have its own rule in order to obtain subjective trust values matching their requirements. For instance, in contrast from the URLLC example rule, the rule base of an MVNO offering only *eMBB* might be less stringent on packet drop but requires *good* or *excellent* bandwidth in addition to latency. In the proposed

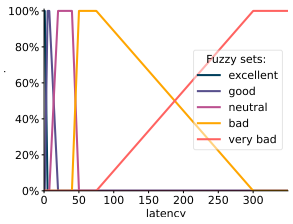

(a) Latency membership function for the complete latency definition space.

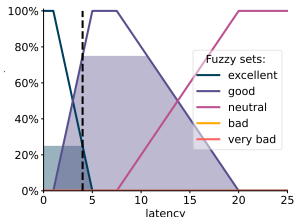

(b) Latency memberships function zoomed for 0-25 ms.

Fig. 2: *Fuzzy membership function associates a crisp value (here, the latency defined in ms) to its corresponding set of terms: excellent, good,... On Figure 2b for a latency crisp value of 4ms (materialized by the dashed vertical black line) the resulting fuzzy set is represented by the colored areas, here the result is 25% excellent and 75% good.*

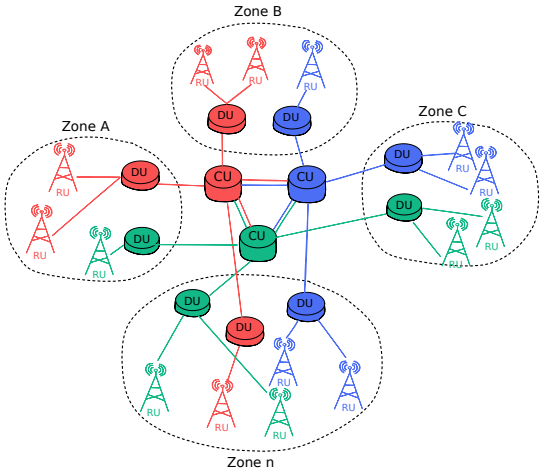

Fig. 3: *Schematic representation of the considered ran sharing topology. In this representation 3 Mobile Network Operator (MNO)s have a partial coverage on the 4 different geographical zones, they must lease DU and RU resources from each others to have a complete coverage for their different clients.*

approach, the shared feedback $f^n_{i,k}$ is a vector containing the crisp values obtained by monitoring the result of a transaction. Each MVNO then creates its own rule base to accommodate its use case. Leveraging this fuzzy logic process each MVNO can obtain a subjective reputation score, enabling him to choose the available InP best suited to his need.

IV. Evaluation use case

This section details the simulation choices that have been made to test our proposed reputation system.

A. Participants typology

We simulate a general O-RAN sharing scenario with 15 participants that are both MVNO and InP, they are diverse in terms of:

1) *geographical coverage*, as illustrated in Figure 3 and detailed in Table II. We create ten different zones so that none of the InPs have a total coverage. 8 different participants are present on each zone. Similar to Zeydan *et al.* [2] *geographical split*, MVNOs must lease NS to InPs on a geographical basis to extend their coverage,
2) *capabilities*, we consider here that each InP provides a unique use case, and each MVNO consumes a unique use case as well. For instance, an InP providing eMBB services will perform poorly on URLLC requests. For each zone, the capacity per use case thus depends on the capabilities of the InPs that are present. The distribution of these capacities is presented in Figure 4.

TABLE II: Evaluation parameters

Simulation parameters	
Slices types	eMBB, URLLC, mMTC
Number of MVNOs & InPs	15: 5 eMBB, 5 URLLC, 5 mMTC
Number of geographical zones	10
Each MVNO number of interactions	150
Interactions duration	0.025
Simulation duration	15
Inconsistent InPs failure rate β	0.05
Reputation and decay parameters	
$\overrightarrow{fs_{i,k}}$ length	0.1
$\overrightarrow{fl_{i,k}}$ length	5
λ	0.5
ε	0.1

For more information on the necessary advertising, biding and monitoring interfaces to build a multi-tenant O-RAN marketplace, refer to [13].

B. Considered topology

As illustrated in Figure 3, participants share their RU and DU but not their CU. There are two main reasons for this choice.

1) The latency constraints are much higher for RU↔DU link (typically 1 or 10 ms) than for the DU↔CU link, thus, the CU can be on a more remote location.
2) The CU hosts the Service Data Adaptation Protocol (SDAP) and Packet Data Convergence Protocol (PDCP) layers in the O-RAN functional split. SDAP layer is in charge of associating QoS flow with data flow, it therefore has the knowledge to prioritize flows if necessary. The PDCP layer is notably in charge of the ciphering and integrity of the transmitted data, leaving it to MVNO ensure that the InP cannot tamper with nor access the data he transmits.

C. Simulation scenarios

In order to evaluate the proposed reputation system, we test different behavior for the InPs. First, we consider a `base` behavior where the InP properly provides the requested service. Secondly, as we want to test the ability of the tested decay function to react to a punctual network failure (**constraint 1**), we introduce `outage` participants that experience a punctual and sudden drop in service quality but behave normally otherwise. The outage period is illustrated in Figure 5 with

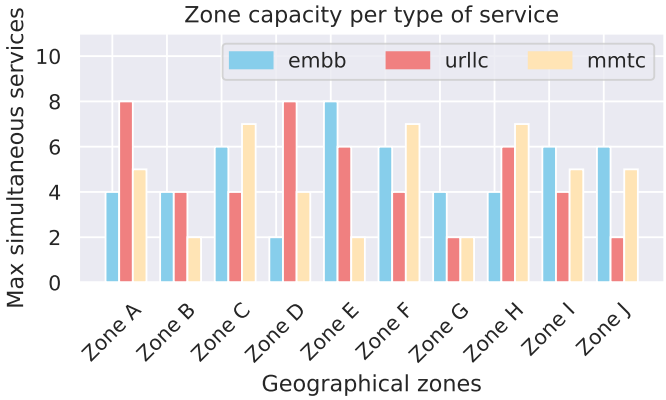

Fig. 4: *Maximum number of simultaneous transactions on each geographical zone per slice types.*

the red top line. Finally, as we want to prevent *inconsistent behaviors* we introduce `oscillating` InPs that provide *normal* services most of the time but degrade their services at a β rate. For the `outage` and `oscillating` scenario, two of the five participants from each use case modify their behavior, leaving three participants per use case with an unchanged behavior.

D. Implementation

The code for all experiments can be found online[1], with configuration and seeds for each considered baseline and evaluation scenario. We also provide poetry lock files to enable anyone to reuse the same software versions as in this paper. To reduce the bias induced by drafting a single seed, all tests were carried out on 30 different seeds, the results presented are an average of these seeds.

V. Evaluation

We want to show that the proposed reputation system can redirect MVNOs to InP that match their requirements, while limiting recommendations to InP that demonstrate *inconsistent behavior*. This can be done by monitoring the number of negative interactions, defined as interactions that terminate with a tenant MVNO declaring to be either *very unsatisfied*, *unsatisfied* or *neutral*.

A. Matching MVNOs with adequate InPs

In Figure 5, we present the number of negative interactions overtime for the different scenarios. Specifically, Figure 5a, depicts a base scenario where all InPs are trustworthy and reliable. The red starred line materializes the naive approach in which MVNOs doesn't rely on reputation but randomly select InPs. This approach maintains a high number of failures throughout the simulation. The other line shows different reputation-based approaches, they also start with a fairly high number of negative interactions. This can be explained by the fact that the MVNOs have no information about the type of services provided by the InPs. An adjustment phase is therefore necessary to find out which InPs are likely to meet their needs. Once this initial exploration phase is over, the number of negative interactions falls sharply and remains low for the rest of the simulation. We can therefore conclude that, at the cost of an initial exploration, the reputation system in place enables us to redirect MVNOs towards InPs corresponding to their needs when the InPs behave as expected.

B. Reacting to inconsistent behaviors

1) `Outage` *scenario:* in this scenario, some InPs are subject to an outage between times 5 and 7 of the simulation. On Figure 5b which represents the total number of negative interactions, we can observe a peak at the beginning of the outage. The number of negative interactions then decreases as the reputation system adjusts its score to the outage. This reputation adjustment is also visible on Figure 8 which underlines participant's reputation going down as they are going through the outage.

Another point, which is apparent on Figure 8 is that when the outage ends, the reputation of the involved InPs doesn't increase back. A direct effect of this permanent reputation drop is the reduction of interactions for these InPs. This effect can be seen in Figure 6, the drop in total interaction continues at then end of the simulation even while the outage has been long finished. In summary, the two decay respect **constraint 1** but fail for **constraint 2**.

2) `Oscillatory` *scenario:* this scenario is mostly exposed in Figure 7. First, we can observe that the reputation system limits the impact of this *inconsistent behavior*: the InPs with oscillatory behaviors have less transaction opportunity than the one with the base behavior. Additionally, and as expected from Section III-A, due to its longer memory span, the adaptive decay is slightly more efficient than the exponential decay in detecting `oscillatory` InPs and limiting their number of interactions. This is especially apparent for the mMTC oscillatory InPs who are gaining more interactions towards the end of the simulation in the exponential case but not in the adaptive one. Nonetheless, comparing Figure 7 with Figure 6 we can see that both solutions penalize the oscillating InPs less than the ones that suffered from an outage.

VI. Conclusion

We consider a market in which MVNOs rent resources from InPs that they don't necessarily trust. To solve this trust problem in the context of 5G, where the needs of MVNO may be very different, we propose a reputation system based on fuzzy logic. We evaluate this system using a RAN sharing simulation that takes into account the geographical distribution and specialization of the various players. The simulation shows that, after an initial identification phase, the system is able to redirect each MVNO to an InP suited to its need. We test two decay functions from the literature, both to limit the effects of a one-off failure or of an InP behaving inconsistently. However, using these functions, an InP suffering from a punctual outage will be durably affected. In the future, we plan to solve this problem with a new decay function.

[1] https://github.com/lekda/5G-fuzzy-reputation

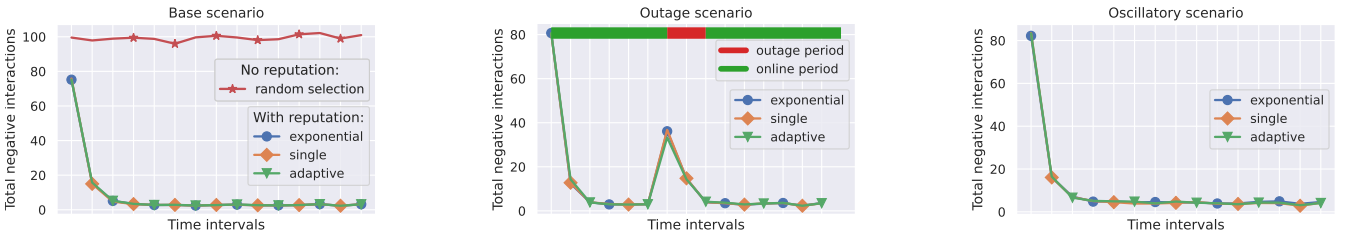

(a) *The starred red line show the number of negative interactions when no reputation is in place. Other lines represent the reputation system results on different decay strategies.*

(b) *During the outage period (materialized in red) the total number of failures temporary increases and then quickly reduces for all tested decay functions.*

(c) *The total number of failures in the oscillatory scenario is comparable to the base scenario for all tested decay functions.*

Fig. 5: *Number of negative interactions overtime on the 15 time intervals. Each sub-figure represents one different behavior, the lines are a comparison of the different decay strategies tested on the reputation system.*

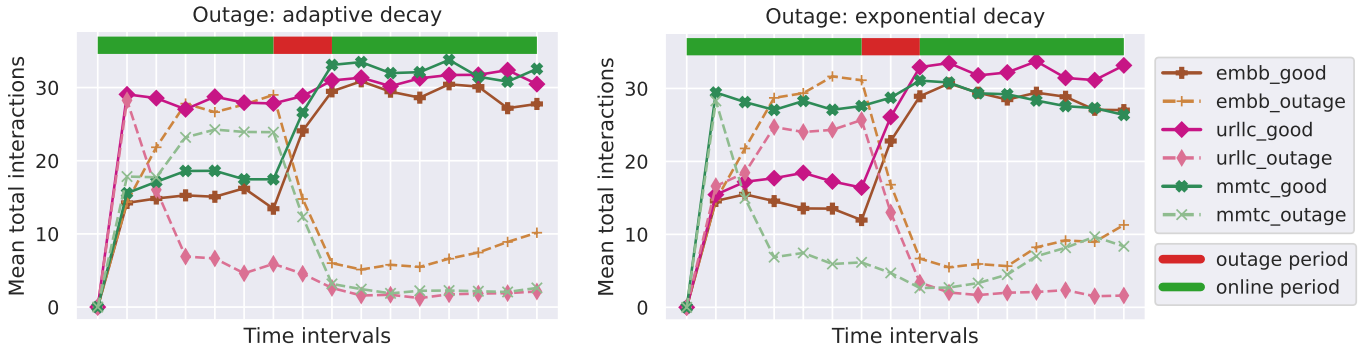

Fig. 6: *Mean total number of interactions (both positives and negatives one are included) for each category of participants. For both tested decay functions, InPs that experience an outage see a sustained and significant reduction of transactions.*

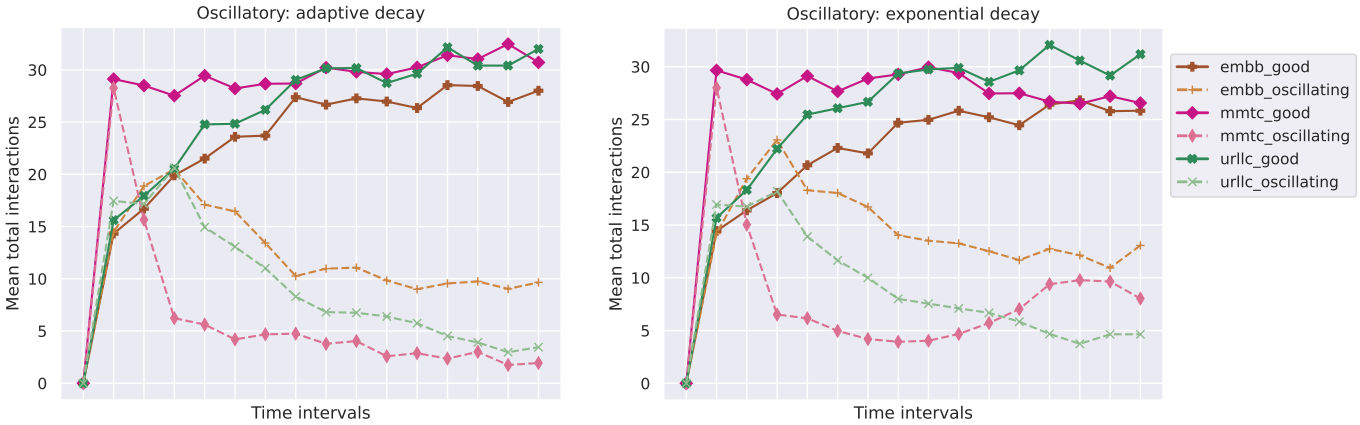

Fig. 7: *Mean total number of interactions (both positives and negatives one are included) for each category of participants. The oscillating participants have a small but continuous failure rate ($\beta = 0.05$). Both decay functions reduce the number of interactions on oscillating InPs. As expected, this effect is slightly more pronounced on the adaptive window decay solution. Furthermore, both solutions penalize the oscillating InPs less than the ones that suffered from an outage seen in Figure 6.*

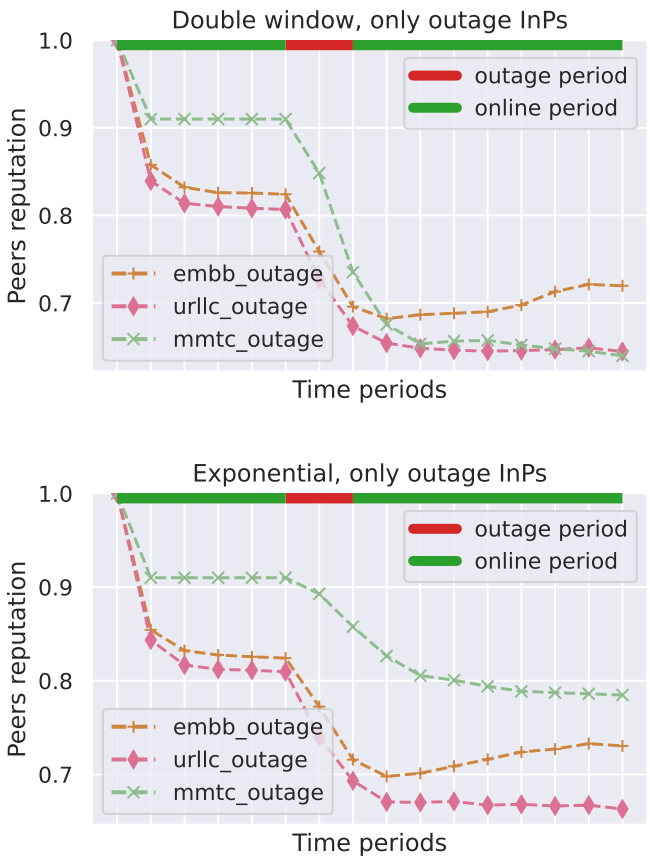

Fig. 8: *Mean reputation of the participants suffering from an outage. For the sake of readability, only the reputation from participants with similar use case is used on the average.*

References

[1] N. Slamnik-Kriještorac, H. Kremo, M. Ruffini, and J. M. Marquez-Barja, Sharing Distributed and Heterogeneous Resources toward End-to-End 5G Networks: A Comprehensive Survey and a Taxonomy. *IEEE Communications Surveys & Tutorials*, 2020.

[2] E. Zeydan, S. S. Arslan, and Y. Turk, Exploring Blockchain Architectures for Network Sharing: Advantages, Limitations, and Suitability. *IEEE Transactions on Network and Service Management*, 2023.

[3] M. Kassis, S. Costanzo, and M. Yassin, Flexible Multi-Operator RAN Sharing: Experimentation and Validation Using Open Source 4G/5G Prototype., Jun. 2021.

[4] R. Ou, G. Sun, D. Ayepah-Mensah, G. O. Boateng, and G. Liu, Two-Tier Resource Allocation for Multitenant Network Slicing: A Federated Deep Reinforcement Learning Approach. *IEEE Internet of Things Journal*, Nov. 2023.

[5] G. Sun, D. Ayepah-Mensah, G. O. Boateng, N. A. Kuadey, M. B. Omer, and G. Liu, Holistic Roadmap of Trends in Radio Access Network Slicing: A Survey. *IEEE Communications Magazine*, Dec. 2023.

[6] G. O. Boateng, D. Ayepah-Mensah, D. M. Doe, A. Mohammed, G. Sun, and G. Liu, Blockchain-Enabled Resource Trading and Deep Reinforcement Learning-Based Autonomous RAN Slicing in 5G. *IEEE Transactions on Network and Service Management*, Mar. 2022.

[7] J. Pérez-Romero, O. Sallent, R. Ferrús, and R. Agustí, Profit-Based Radio Access Network Slicing for Multi-tenant 5G Networks., Jun. 2019.

[8] I. Vilà, J. Pérez-Romero, O. Sallent, and A. Umbert, A Multi-Agent Reinforcement Learning Approach for Capacity Sharing in Multi-Tenant Scenarios. *IEEE Transactions on Vehicular Technology*, Sep. 2021.

[9] T. Faisal, M. Dohler, S. Mangiante, and D. R. Lopez, BEAT: Blockchain-Enabled Accountable and Transparent Network Sharing in 6G. *IEEE Communications Magazine*, Apr. 2022.

[10] P. Resnick, K. Kuwabara, R. Zeckhauser, and E. Friedman, Reputation systems. *Communications of the ACM*, 2000.

[11] G. He, W. Su, S. Gao, N. Liu, and S. K. Das, NetChain: A Blockchain-Enabled Privacy-Preserving Multi-Domain Network Slice Orchestration Architecture. *IEEE Transactions on Network and Service Management*, Mar. 2022.

[12] J. M. J. Valero, V. Theodorou, M. G. Pérez, and G. M. Pérez, SLA-driven trust and reputation management framework for 5G distributed service marketplaces. *IEEE Transactions on Dependable and Secure Computing*, 2023.

[13] L. Giupponi and F. Wilhelmi, Blockchain-enabled Network Sharing for O-RAN in 5G and Beyond. Dec. 2021.

[14] E. Koutrouli and A. Tsalgatidou, Taxonomy of attacks and defense mechanisms in P2P reputation systems—Lessons for reputation system designers. *Computer Science Review*, May 2012.

[15] L. Xiong and L. Liu, PeerTrust: Supporting Reputation-Based Trust for Peer-to-Peer Electronic Communities. *IEEE Transactions on Knowledge and Data Engineering*, Jul. 2004.

[16] R. Figueiredo, S. Eisa, and M. L. Pardal, SureRepute: Reputation System for Crowdsourced Location Witnesses., Dec. 2022.

[17] A. Jøsang and R. Ismail, The Beta Reputation System. *In: Proceedings of the 15th Bled Conference on Electronic Commerce*, Jan. 2002.

[18] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina, The Eigentrust algorithm for reputation management in P2P networks., ser. WWW '03, New York, NY, USA: Association for Computing Machinery, May 2003.

[19] A. Almogren, I. Mohiuddin, I. U. Din, H. Almajed, and N. Guizani, FTM-IoMT: Fuzzy-Based Trust Management for Preventing Sybil Attacks in Internet of Medical Things. *IEEE Internet of Things Journal*, Mar. 2021.