



HAL
open science

A decentralised data layer for collaborative End-to-End service assurance

Vincent Messié, Benoit Radier, Veronica Quintuna Rodriguez, Gaël Fromentoux, Sandrine Vaton, Isabel Amigo

► To cite this version:

Vincent Messié, Benoit Radier, Veronica Quintuna Rodriguez, Gaël Fromentoux, Sandrine Vaton, et al.. A decentralised data layer for collaborative End-to-End service assurance. ICIN 2022: 25th Conference on Innovation in Clouds, Internet and Networks, Mar 2022, Paris (virtual conference), France. 10.1109/ICIN53892.2022.9758094 . hal-03647594

HAL Id: hal-03647594

<https://imt-atlantique.hal.science/hal-03647594>

Submitted on 20 Apr 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A decentralised data layer for collaborative End-to-End service assurance

Vincent Messié * [†], Benoit Radier *, Veronica Quintuna Rodriguez *,
Gaël Fromentoux *, Sandrine Vaton [†], Isabel Amigo [†]

* Orange Labs, 2 Avenue Pierre Marzin, 22300 Lannion, France
{vincent.messie, benoit.radier, veronica.quintunarodriguez, gael.fromentoux}@orange.com

[†] IMT Atlantique, 655 Avenue du Technopôle, 29280 Plouzané, France
{vincent.messie, sandrine.vaton, isabel.amigo}@imt-atlantique.fr

Abstract—This paper proposes a solution for validating an End-to-End service chain built by multiple actors that may not trust each other. We notably introduce a “data layer” powered by a Distributed Ledger (DL, a.k.a “Blockchain”) using a Directed Acyclic Graph (DAG). This component will enable all players involved in a network service chain to share trusted and secure performance data, whilst avoiding the participation of trusted third parties. We consider as a driving use case a scenario where resource providers and resource consumers (a.k.a “prosumers”) interact together to build on-demand network services. We thus focus on a Cloud-based Radio Access Network scenario and anticipate network disaggregation, allowing the infrastructure to be shared between multiple providers. We show through simulation that the usage of a DAG-based ledger will make the proposed data layer scalable despite the amount of performance data required for monitoring.

Keywords: Blockchain, service chaining, resource sharing, marketplace, DLT.

I. INTRODUCTION

The evolution of telecommunication networks towards cloud-native environments in addition to the diversification of customer needs has given rise to a deep transformation in network ecosystems up to so-called “disaggregated” networks. From a business perspective, network evolution has enabled the emergence of new markets involving various telecom assets providers of different nature (Towercos, Multi-access Edge Computing, Cloud, and connectivity). This opens the door to network marketplace platforms that put into relation *resource (asset) providers* and *resource (asset) consumers (prosumers)* within Business-to-Business models. When various players are involved in a single End-to-End (E2E) service, having trusted, non-repudiable agreements among the parties, it is compulsory to assure the required Quality of Service (QoS) for end-users. New trustfulness technologies such as the DL Technology (DLT) allows the marketplace to become decentralised, thus simplifying asset exchanges [1], [2], [3]. However, a trusted marketplace is not enough to assure E2E services, as it is also required for prosumers to guarantee negotiated Service Level Agreements (SLAs) [1].

Various studies show a growing interest in DLs to foster trust on decentralised systems. Firstly, various initiatives led by the TMForum, including Proofs of Concept [1] and technical reports [2] explore the benefits, opportunities and challenges of DL-based collaborative networks. There are also various use-cases emerging in the network area, such as BALAdIN [4], [5] proposing a novel Blockchain-based network model, and Payflow [6] providing a decentralised solution to automatically trade bandwidth through a DL.

The main contributions of this work are:

- We introduce a “*data layer*” enabling all players involved in a E2E service chain to share performance metrics in a trusted environment, thus achieving a decentralised monitoring;
- We then propose to use the DL technology as the core technology of the data layer to avoid costly trusted third parties. We select more particularly the Tangle, a DL based on DAGs designed for efficiency [7], [8], [9].
- And we finally show through simulation that the Tangle can scale and sustain the exposed data layer’s use case.

This paper is organised as follows. In section II, we describe the DLs, while focusing on DAG-based ledgers. In section III, we describe the driving scenario where providers share part of their infrastructure to allow consumers to build federated E2E service chains. We particularly describe the proposed DL-based data layer used to validate SLAs in a decentralised way. In section IV, we conduct simulations of the data layer using the Tangle. We introduce a “sampled” model based on Popov’s works [7] and enriched by the Bramas model [10]. Results confirm the pertinence of the Tangle for the proposed use-case. Conclusions and next steps are presented in section V.

II. GETTING TRUST BY DISTRIBUTED LEDGER TECHNOLOGY

A. On Distributed Ledger Technology

Democratised in 2008 thanks to the Bitcoin currency [11], the DL technology adds trust to any distributed

TABLE I
DAG-BASED DL TECHNOLOGIES

	Nano [13]	Hashgraph [14]	IoTa (1.0) [7]
Launch year	2018	2018	2016
Project openness	Open-source	closed-source (open-review)	Open-source
Governance	Decentralised	Fixed number of validating nodes	Central entity (Coordinator)
Smart-contracts	No	Yes	Yes

applications. A DL is a database distributed on a network of *participating nodes* in a trustless environment. The ledger is replicated on each participating node, and new data is broadcasted on the network, to be then approved by every participant. This database is “add-only”, which means stored data (a.k.a “transactions”) cannot be altered, thanks to cryptography-based mechanisms, e.g. Proof of Work (PoW) [11], [12]. These mechanisms allow the DL to reach consensus, while preventing malicious nodes to corrupt it.

The most widely used DL technology is the *Blockchain*: the incoming data are packed into blocks, that are then chained together so that each block validates its predecessor [11], [12]. However, the main issue of the Blockchain is its synchronous nature making it inefficient, restraining it to scale well, and implying a significant delay for transaction validation.

Various studies address such issues as being the main drawbacks of nowadays DLs [8], [9].

Major protocol changes are proposed to overcome these issues and achieve scalability, such as DAG-based solutions like Nano [13], Hashgraph [14], or IoTa [7].

B. On Directed Acyclic Graph based DLs

Table I lists some existing DAG-based DL technologies, and their key properties. It shows that the Tangle is the oldest, most open and flexible DAG technology we know of. Furthermore, it is also widely used at this time [7], [8], [9], [10], [15]. As a result, we focus on the Tangle in this paper.

The Tangle was designed to efficiently handle micro-transactions. A Tangle can be seen as a “two-dimensional Blockchain”.

With the Tangle, transactions are not packed into blocks but rather work in a standalone mode to form a DAG. To be attached to the DAG, any transaction needs to select and approve two of its non-validated predecessors (referred to as “tips” of the Tangle). To select their parents, nodes implement a Tip Selection Algorithm (TSA).

As shown in Figure 1, the highlighted red transaction indirectly validates dashed blue transactions, while being validated by dashed orange ones. The solid blue transactions represent “tips” (unvalidated transactions).




Fig. 1. A Tangle-based ledger

While a Blockchain like Bitcoin relies on a PoW or other consensus mechanism to achieve trust, the key of the Tangle technology lies in its TSA.

The nature of the Tangle thus makes transaction commitment and validation more challenging. In current implementation, a “coordinator” is implemented for securing transaction commitment [16]. However, the Tangle asynchronous nature allows more scalability and efficiency than a Blockchain. The transaction processing capacity of a Tangle is then theoretically unbounded, and the network better scales with a growing number of participants.

A simplified mathematical model of the Tangle is introduced in [7]. It shows that the Tangle remains stable over time, and may reach consensus in a stochastic way. Bramas [10] has then proposed a further simplified model, demonstrating the stability through a formal analysis.

In the following sections we consider a Tangle-based DAG as the driving DL for the proposed data layer.

III. MULTI-ACTOR SERVICE CHAINING MODEL

A. Multi-actor E2E service chaining scenario

With the emergence of virtualisation, network operators are migrating from dedicated physical to Virtualised Network Functions (VNFs) running on banalised hardware.

As a driving scenario we consider a shared environment hosting mobile network elements.

We particularly consider a Cloud Radio Access Network (RAN) architecture, disaggregating the RAN functions in antennas, Edge and central Clouds [17], [18]. Each of these elements are interconnected by optical networks.

Thus, the hosting physical infrastructure can be shared between multiple providers, enabling telcos to create value while reducing capital expenditures [19]. Also, a given cloud infrastructure may give host to multiple E2E services.

As an example presented figure 2, we consider that a given mobile operator (in red) aiming to deploy a mobile network uses (i) a Tower, (ii) an Edge environment close to the chosen antennas, and (iii) a central Cloud infrastructure to deploy its required VNFs. Similarly, another mobile operator (in orange) deploys his virtual network eventually on the same physical infrastructure. Both E2E




Fig. 2. Resource sharing in convergent mobile networks (R/D/CU = Radio/Decentralised/Centralised Unit). Each coloured rectangle represents an actor sharing part of his infrastructure.

network services share the same optical networks, eventually provided by distinct providers.

Recent works have investigated the main challenges of cloud architectures when dealing with latency-sensitive network functions as those of RAN [20], [21], and the need for QoS to be enforced. It is then required to make prosumers pro-actively monitor the service [22], to enforce the QoS. Performance metrics, either customised or standardised [22], [23] and their processing method to monitor the operation of the agreed E2E service, as well as the agents used to collect metrics (referred to as “probes”), and their location in the network are identified on the service creation.

B. Proposal of a DL-based data layer

We then introduce a DL-powered trusted data layer to collect and validate trusted performance metrics. This approach is more efficient than a centralised one, as it doesn’t require an external third party to secure the performance metrics [1], [2].

The proposed architecture is illustrated in Figure 3. Firstly some specific “Usage Report (UR) agents” are deployed as close as possible to the VNFs to collect raw measurements from probes, pack them into transactions, and locally process them by checking their conformity, with the help of applications running on the DL (steps 1 & 2). These metrics come for various locations on the network. Example of metrics are network round trip time, amount of bytes exchanged on a specific point, and also CPU usage of a given VNF [22].

The pre-processed metrics are then stored onto the distributed ledger, and static “full nodes” with higher computing capabilities, hosted by the providers, read them to produce certified E2E KPIs (step 3). Full nodes implement byzantine fault tolerant mechanisms to calculate and certify the KPIs in a decentralised way. By doing so, they validate their commitment, in a similar way than the Iota coordinator. The proposed architecture then allows the production of secured, traceable KPIs that can be used for accounting, reconciliation or SLA enforcement, without relying on any external trusted third party.

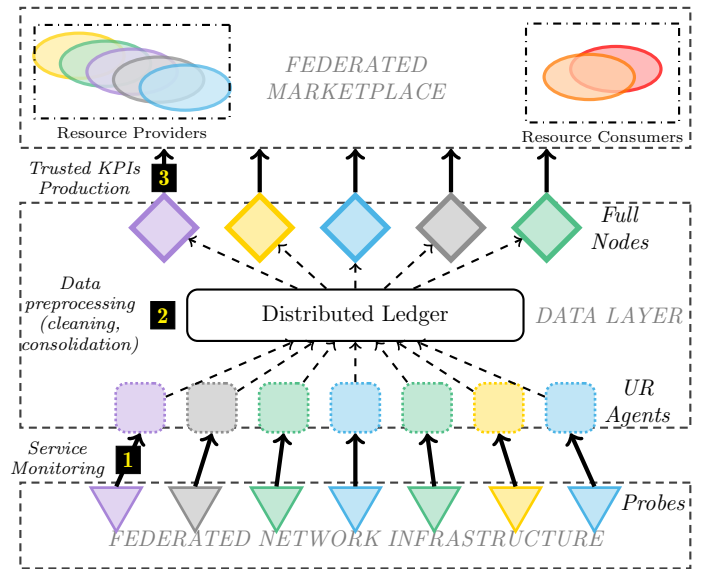


Fig. 3. Data layer architecture (same colours as Fig. 2 for providers and consumers)

C. Assessing the expected transaction throughput

Since SLA enforcement requires near real-time behaviour due to high QoS constraints [20], [21], the main challenge of the proposed data layer is the data rate it can handle, a.k.a the resulting transaction throughput. To roughly compute it, we consider as a numerical example, each active 4G site as an active service chain continuously emitting performance report [22], [23]. There are now 48,500 active 4G sites in France [24].

While already used in nowadays clouds/networks, the exact volume and nature of performance metrics are never disclosed for strategic reasons. In this paper, we consider as a rough estimation that prosumers may need to collect 20-30 raw measurements every 15 minutes for a given E2E service. We then obtain an estimated rate of Transactions per Second (TPS) of $48,500 \times 30 \times \frac{1}{900} \approx 1617tx.s^{-1}$. This number is expected to fluctuate, as the marketplace is dynamic and E2E services can be deployed and terminated on-the-fly. Furthermore, any prosumer may join or leave the marketplace at any time.

A similar study shows the difficulty of such an use case while using a Blockchain [5]. The authors particularly address the need of sharding (i.e splitting) the Blockchain to avoid congestion and enable scalability. However, sharding is not necessary when using a DAG-based DL since it naturally scales. Furthermore, the technology have enough flexibility to accommodate for the TPS fluctuation. In the following section we propose an implementation of the data layer with the Tangle, and we assess its behaviour through simulations.

IV. EXPERIMENTATION

We address in this section the implementation of the data layer Tangle-based DL, and we create a simulator to assess the behaviour of the Tangle when confronted to a high load of transactions. We use the input parameters estimated in the previous section (48,500 4G sites, 20–30 raw measurements every 15 minutes).

A. On Tangle Modelling

The model introduced by [7] represents the DAG as a continuous time stochastic process. The flow of arriving transactions is modelled as a Poisson process of rate λ . Using a simple, random uniform TSA, each arriving transaction randomly attaches itself to two unconfirmed ones (a.k.a the “tips” of the Tangle).

Let $\mathcal{T}(t)$ be the set of tips of the Tangle at time t .

Let h be the time necessary for attaching a transaction to the Tangle, so that a transaction created and visible at time t performed the TSA on $\mathcal{T}(t-h)$. According to Popov [7], one may expect the number of tips $L(t) = \#\mathcal{T}(t)$ to be stationary, fluctuating around a constant value $L_0 = 2\lambda h$. As a result $L(t)$ does not tend to infinity, making the graph stable. Upon attachment of a given transaction, one may expect every new arriving transaction to indirectly validate it after a certain amount of time called the “adaptation period”, thus making the Tangle reach consensus. [7] has estimated this adaptation period as follows:

$$t_0 \lesssim 2.84 \times h \ln L_0 \quad (1)$$

B. Tangle simulator

We then create a simulator¹ of the growth of the Tangle to test the use case. The program generates dummy transactions (nodes of the graph) that are linked (edges of the graph) to two parents. Apart from its two parents, the primary attribute of a transaction is its time of creation, represented as an integer.

The time is discretised to time slots referred to as “rounds”, separated by an interval of Δt . For each round, a random variable following a Poisson law of parameter $\lambda \cdot \Delta t$ is calculated.

We choose $\lambda \cdot \Delta t \ll 1$, so that the poisson law can be approximated by a Bernoulli variable of parameter $\lambda \cdot \Delta t$. If the outcome of the variable is positive a single transaction is then created at the current round.

Any new created transaction will then attach itself to two parents, chosen randomly in the list of transactions that were unvalidated (tips) $h/\Delta t$ rounds ago. As a result, Δ_t is chosen so that $h/\Delta t \in \mathbb{N}$.

The global state of the graph is managed by a “Tangle” object, whose main purpose is to keep track of the unconfirmed transactions (“tips”), and the evolution of the Tangle over time.

¹Simulator is in open source at <https://gitlab.com/vmessie/dag-simulator>

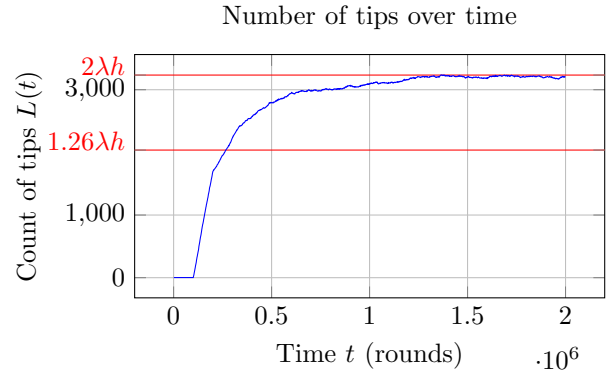


Fig. 4. $L(t)$ with $\lambda = 1617tx.s^{-1}$, $h = 1s$, $\Delta t = 10^{-5}s$

Two “genesis” transactions are created at $t = 0$, as the initial state of the graph. Then a loop iterates on t until a given number of rounds (time frame).

When a transaction is created at time t , the list of tips $\mathcal{T}(t)$ is saved on a sorted list. This list can be then looked up on the creation of future transactions, to account for the $h \cdot \Delta t$ value, thus allowing new transactions to select their tips in the past. As this list is already sorted, the complexity of the lookup grows linearly with the amount of created transactions.

This list is then also used at the end of the simulation, to extract the evolution of the number of tips over time.

We run the simulation with $\lambda = 1617tx.s^{-1}$ as the value we have previously estimated in section III-C.

Let us now estimate the h value, a.k.a the delay on attaching a transaction to the graph. Work cited in [15] has shown that as few as 687 ms are needed to prepare and sign a transaction on embedded systems without performing a PoW. We will then take $h = 1s$, as the remaining 313 milliseconds would then reflect residual time to compute the transaction’s payload, and network propagation. We define $\Delta t = 10^{-5}s$. The simulation is then run on a time frame of 20 seconds or $2 \cdot 10^6$ rounds.

C. Simulation results

We get the number of tips over time $L(t)$, to validate the stability of the Tangle in the proposed use-case, and extract the L_0 value.

Figure 4 displays $L(t)$ as produced by the simulator. While starting at 2 (at $t = 0$, there are only the genesis transactions), we observe $L(t)$ to grow, then reach and fluctuate around a constant value, which validates former hypothesis. However, this constant value is closer to $L_0 = 2\lambda h$ than $1.26\lambda h$, which makes the model close to the continuous one [7].

Equation (1) gives a t_0 of 23 seconds, i.e. it will take 23 seconds for a given transaction to be validated by every incoming transaction. When such an event occurs, we also consider that the transaction is then known to every node of the data-layer.

As a result, we can expect the data layer to process and validate KPIs in less than 30 seconds.

D. Discussion & Next Steps

Through these simulations we have shown that the Tangle will be able to scale and sustain the proposed data layer, given the use of a basic, random uniform TSA, and without any PoW mechanism. Indeed, the Tangle remains stable with the input parameters presented sections III-C and IV-B.

To further evaluate the data layer on a realistic environment, next step is to implement the proposed solution in a testbed which will be composed of a Cloud-RAN network, and a private Tangle instance to implement the data layer.

Some IoTa nodes (“UR Agents”) will also be deployed on the network infrastructure itself. It shall enable to further evaluate the overhead and resource consumption induced by the operations of the data layer.

V. CONCLUSION

We have proposed a DL-based data layer to validate SLAs of multi-actor E2E service chains in a trusted and distributed way, by sharing performance metrics between multiple players. By storing performance metrics in a DL, the proposed model ensures their integrity without relying on a trusted third party. We have introduced the required components (namely “Usage Report agents” onboarded on each E2E service, and static “full nodes”) to support the proposed enforcement mechanisms.

We then select the Tangle as the driving DL for the data layer. The complexity of Tangle-based solutions is on their stability when they scale. We particularly validate the scalability of the proposed model by simulation by using a sampled time model.

Simulation results show that a nation-wide Tangle instance remains stable and consistent over time, and shall efficiently sustain the data layer, as long as the time to produce a transaction remains low enough (1 second). In this context, European or worldwide Tangle instances should also reach consensus with full nodes deployed at a national level, thanks to its asynchronous behaviour.

Next step is the experimentation while using a Cloud-RAN network and a private Tangle instance powering the data layer.

REFERENCES

- [1] A. Adhiappan *et al.*, “Federated CSPs marketplace : A DLT-based data trust enabling business assurance for CSPs platforms federation,” TMForum, Tech. Rep., catalyst C20.0.34 White Paper. [Online]. Available: <https://www.tmforum.org/vertical-industry-telcos-federated-dlt-based-marketplace/>
- [2] “CSP use cases utilizing blockchain,” TMForum, Tech. Rep. TR279. [Online]. Available: <https://www.tmforum.org/resources/technical-report/tr279-csp-use-cases-utilizing-blockchain-v3-1/>
- [3] Blockchain unleashed - TM forum. Library Catalog: www.tmforum.org. [Online]. Available: <https://www.tmforum.org/catalogs/blockchain-unleashed/>
- [4] V. Messié, G. Fromentoux, X. Marjou, and N. Labidurie, “Baladin for blockchain-based 5g networks,” in *2019 22nd Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN)*, 2019, pp. 201–205.
- [5] V. Messié, G. Fromentoux, N. Labidurie, B. Radier, S. Vaton, and I. Amigo, “BALAdIN: truthfulness in collaborative access networks with distributed ledgers,” *Annals of Telecommunications*, Jun. 2021. [Online]. Available: <https://doi.org/10.1007/s12243-021-00855-x>
- [6] D. Chen, Z. Zhang, A. Krishnan, and B. Krishnamachari, “Payflow: Micropayments for bandwidth reservations in software defined networks,” in *2019 International Conference on Computer Communications (INFOCOM)*. IEEE, 2019, pp. 26–31.
- [7] S. Popov, “The tangle,” Tech. Rep., 2018, IoTa Whitepaper. [Online]. Available: https://iota.org/IOTA_Whitepaper.pdf
- [8] K. Yeow, A. Gani, R. Ahmad, J. Rodrigues, and K. Ko, “Decentralized consensus for edge-centric internet of things: A review, taxonomy, and research issues,” *IEEE Access*, vol. 6, pp. 1513–1524, 2018.
- [9] Y. Xiao, N. Zhang, W. Lou, and Y. T. Hou, “A survey of distributed consensus protocols for blockchain networks,” *IEEE Communications Surveys & Tutorials*, 2020.
- [10] Q. Bramas, “The Stability and the Security of the Tangle,” in *International Conference on Blockchain Economics, Security and Protocols (Tokenomics 2019)*, ser. OpenAccess Series in Informatics (OASIs), vol. 71, 2020, pp. 8:1–8:15.
- [11] S. Nakamoto *et al.*, “Bitcoin: A peer-to-peer electronic cash system,” Tech. Rep., 2008, bitcoin Project White Paper. [Online]. Available: <https://bitcoin.org/bitcoin.pdf>
- [12] G. Wood *et al.*, “Ethereum: A secure decentralised generalised transaction ledger,” Tech. Rep., 2014, ethereum Project Yellow Paper. [Online]. Available: <http://gavwood.com/Paper.pdf>
- [13] C. LeMahieu, “Nano: A feeless distributed cryptocurrency network,” Tech. Rep. [Online]. Available: <https://nano.org/en/whitepaper>
- [14] D. L. Baird, M. Harmon, and P. Madsen, “Hedera: A Public Hashgraph Network & Governing Council,” Tech. Rep. v2.1, 2020. [Online]. Available: <https://hedera.com/hh-whitepaper>
- [15] D. Stucchi, R. Susella, P. Fragneto, and B. Rossi, “Secure and effective implementation of an IOTA light node using STM32,” in *Proceedings of the 2nd Workshop on Blockchain-enabled Networked Sensor*, ser. BlockSys’19. Association for Computing Machinery, pp. 28–29. [Online]. Available: <https://doi.org/10.1145/3362744.3363344>
- [16] S. Popov *et al.*, “The coordicide,” Tech. Rep., 2020, IoTa Coordicide Whitepaper. [Online]. Available: https://files.iota.org/papers/20200120_Coordicide_WP.pdf
- [17] V. Quintuna Rodriguez and F. Guillemin, “Higher aggregation of gNodeBs in Cloud-RAN architectures via parallel computing,” in *22nd Conference on Innovation in Clouds, Internet and Networks and Workshops, ICIN*, 2019, pp. 151–158.
- [18] ORAN alliance. Library Catalog: www.o-ran.org. [Online]. Available: <https://www.o-ran.org>
- [19] Blockchain-based telecom infrastructure marketplace - TM forum. Library Catalog: www.tmforum.org. [Online]. Available: <https://www.tmforum.org/blockchain-based-telecom-infrastructure-marketplace/>
- [20] V. Quintuna Rodriguez, “New Network / IT Command: Virtualized Function Performance for a Programmable Infrastructure,” Sorbonne University, 2018. tel-01884431f.
- [21] V. Quintuna Rodriguez and F. Guillemin, “Cloud-RAN Modeling Based on Parallel Processing,” *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 3, pp. 457–468, March 2018.
- [22] 3GPP TS 28.554 V16.03.0, “Management and orchestration; Concepts, use cases and requirements (Release 15),” 2019.
- [23] 3GPP TS 28.552 V16.04.0, “Management and orchestration; Concepts, use cases and requirements (Release 15),” 2019.
- [24] Observatoire ANFR : près de 48 500 sites 4g en service en france au 1er octobre. [Online]. Available: <https://www.anfr.fr/toutes-les-actualites/actualites/observatoire-anfr-pres-de-48-500-sites-4g-en-service-en-france-au-1er-octobre/>