

DBNS: A Distributed Blockchain-Enabled Network Slicing Framework for 5G Networks

Mohammed Amine Togou, Ting Bi, Kapal Dev, Kevin McDonnell, Aleksandar Milenovic, Hitesh Tewari, and Gabriel-Miro Muntean

This article introduces a distributed blockchain-enabled network slicing (DBNS) framework that enables service and resource providers to dynamically lease resources to ensure high performance for their end-to-end services.

ABSTRACT

5G technology is expected to enable many innovative applications in different verticals. These applications have heterogeneous performance requirements (e.g., high data rate, low latency, high reliability, and high availability). In order to meet these requirements, 5G networks endorse network flexibility through the deployment of new emerging technologies, mainly network slicing and mobile edge computing. This article introduces a distributed blockchain-enabled network slicing (DBNS) framework that enables service and resource providers to dynamically lease resources to ensure high performance for their end-to-end services. The key component of our framework is the global service provisioning (GSP), which provides admission control for incoming service requests along with dynamic resource assignment by means of a blockchain-based bidding system. The goal is to improve users' experience with diverse services and reduce providers' capital and operational expenditures.

INTRODUCTION

5G technology is expected to promote many new applications in different vertical industries such as manufacturing, automotive, healthcare, energy, and mobile broadband. Such applications will encompass a variety of use cases, each having distinct sets of requirements (e.g., bandwidth, latency, scalability, and availability). A real-life scenario depicting some of these use cases is illustrated in Fig. 1. It consists of sharing high-resolution video feeds along with braking data among vehicles to enable motorists to adjust their driving to traffic conditions and to anticipate sudden stops along with streaming high-definition multimedia content for passengers.

Provisioning resources to fulfill the variety of requirements of these applications is an arduous task given today's *one-size-fits-all* networks. Indeed, accommodating these applications while supporting existing services requires a programmable and flexible network infrastructure that can be tailored to the specific needs of each application. Besides, guaranteeing service continuity between realms belonging to different network providers is still a challenge. For instance, SP1 in Fig. 1 cannot fulfill the blue car's service requests throughout its journey (i.e., from *home* to *office* passing by

the school) as it lacks coverage and computational capabilities. A possible way to alleviate this problem is network sharing. It enables network operators to share their infrastructure to reduce capital/operating expenditures (CAPEX/OPEX) and to offer lower-priced services to customers. However, the way mobile networks are built and operated limits the network operators' ability to create novel services that can support the diverse requirements of emerging 5G use cases [1].

5G integrates two key technologies: network slicing and mobile edge computing (MEC). The former splits the physical network infrastructure into multiple logical networks, called slices, which are controlled and managed independently by slice owners. The latter deploys computing resources near end users to enable low-latency and high-bandwidth network access. While flexibility is emphasized through the network slicing dynamics, providing real-time end-to-end services that can involve different operators is still a challenge. Indeed, there is a clear urgency to design innovative techniques to enable network slicing interoperations that can span over multiple administrative domains. In this regard, a centralized capacity broker was proposed in [2], which is built on top of the Third Generation Partnership Project (3GPP) network sharing management architecture to enable network operators to lease resources on the fly for a particular time period. This scheme was largely improved in [3] to optimize the on-demand allocation of resources based on mobile traffic forecast. An Internet of Things (IoT) broker was proposed in [4] to enable operators to offer network slices as a service in order to optimize slice resource orchestration. In [5], a management and orchestration architecture that deploys a brokering layer based on software defined networking was proposed to create end-to-end slices and allocate computing and storage capacities at three levels: sub-domain, domain, and multi-domain. Finally, a brokering system that uses a reinforcement learning algorithm was introduced in [6] to dispatch resource requests to different network operators considering constraints such as delay and geographical location.

Despite their ingenuity, these approaches along with other existing ones face two main challenges: scalability: as the number of operators (and therefore the number of slices) increases, the broker may be overwhelmed, impacting the

Mohammed Amine Togou, Ting Bi, and Gabriel-Miro Muntean (corresponding author) are with Dublin City University; Kapal Dev and Hitesh Tewari (corresponding author) are with Trinity College Dublin; Kevin McDonnell and Aleksandar Milenovic (corresponding author) are with Huawei Technologies Co. Mohammed Amine Togou and Ting Bi are joint first authors.

performance of the whole ecosystem; and complexity: to use network slices of other operators, memoranda of understanding (MoUs) need to be established based on pre-defined service level agreement (SLA) requirements. This is a lengthy process.

The main contribution of this article is the design of a signaling-based distributed on-demand framework, called distributed blockchain-enabled network slicing (DBNS), to enable end-to-end and on-the-fly leasing of resources to support service continuity while meeting quality of service (QoS) requirements. The key enabler of our distributed on-demand architecture is global service provisioning (GSP), a control and monitoring component that provides admission control for incoming requests along with resource assignment by means of a blockchain-based bidding system.

The rest of this article is organized as follows. First, we describe the technology enablers of DBNS. Then we present the fundamentals of network slicing orchestration and its business requirements. Next, we introduce the DBNS architecture and describe how it can be deployed and used. Afterward, we analyze DBNS performance through a case study that focuses on video streaming applications. Finally, we present our conclusions along with future research directions.

DBNS TECHNOLOGY ENABLERS

After the early deployments of 3G networks, network sharing was introduced to help mobile operators accelerate their network rollouts as well as offer various services to their customers with reduced costs. Passive sharing and network roaming were the initial network sharing solutions. The former designates the sharing of site locations or physically supporting infrastructure of radio equipment (i.e., mast), while the latter denotes the ability of subscribers to use networks of other operators based on contractual agreements [2]. Then active radio access network (RAN) sharing followed. It consists of sharing active network equipment including base stations, antennas, and mobile backhaul equipment [7]. This allowed mobile operators to merge spectrum resources on the basis of contractual agreements. However, with the continuous growth of data traffic generated by a plethora of emerging applications, new challenges that go beyond the original RAN sharing concept have been introduced. As a result, mobile operators had to think of new ways to redesign their networks to address these concerns.

Network slicing is a concept that has been proposed to cope with the surge in mobile data traffic. It is based on two technologies, software-defined networking [8] and network function virtualization, to allow for efficient network management while enabling flexible network customization to meet traffic requirements. Using network slicing, operators can build multiple isolated virtual networks on a shared physical infrastructure. These logical networks, called slices, are self-contained and orchestrated in different ways depending on their service requirements and can be owned by one or multiple tenants. This provides new revenue opportunities for mobile operators as slices can be used by other operators, offering efficient utilization of the infrastructure. Still, innovative approaches for efficient




Figure 1. Sample of real-life use cases for the automotive vertical. The blue car is an autonomous vehicle that requires surrounding cars to share two elements: high-resolution real-time video feeds to adjust the driving to traffic conditions (see-through), and braking data to anticipate sudden stops. Furthermore, high-quality entertainment services (e.g., games and movies) are streamed via its built-in entertainment system to its passengers.

and secure end-to-end network slicing admission control that can span over multiple administrative domains need to be investigated.

Lately, blockchain (BC) has been integrated in 5G networks to mitigate various security concerns as it offers immutability, transparency, decentralization, and privacy. Its applications in 5G can be classified into three categories: supporting 5G technology enablers such as MEC [9]; providing efficient 5G services like network management and virtualization [5]; and endorsing 5G IoT applications [10]. BC is a peer-to-peer decentralized distributed ledger that permanently and chronologically records and ensures immutable and unchangeable truthful transactions in untrusted or partially trusted networks via cryptography-based consensus mechanisms [11]. A simple BC operation goes through the steps illustrated in Fig. 2. First, an entity (e.g., computer) requests a transaction that is broadcasted to all nodes in the peer-to-peer (P2P) network. All nodes in the network validate the transaction using the requested entity's public key. Once verified, the transaction is linked to other transactions to create a new data block, which is appended to the existing blockchain, indicating the completion of the transaction. Once the transaction has been locked into a block, the other participants in the P2P network can create their own transactions and broadcast them over the BC P2P network to be included in the next block on the ledger.

There are two types of BC: public and private. Our focus in this article is on private BC as we aim to design a network where only a finite number of entities can access it [12]. There are multiple platforms that enable the creation of such networks. For instance, Multichain is a permissioned blockchain that allows multiple networks to simultaneously be on a single server and has compatibility with the bitcoin network, including transaction/block formats, P2P protocols, digital signature schemes, and application programming interfaces (APIs). But the main problem with Multichain is the lack of support for smart contracts. Unlike Multichain, Ethereum supports smart contracts; however, it suffers two major problems: anyone can connect to the network, and all data

Given that the network infrastructure may belong to either MNOs or MVNOs, addressing the following questions is of paramount importance for these two players to decide whether or not to invest in network slicing: Is it beneficial to their businesses?

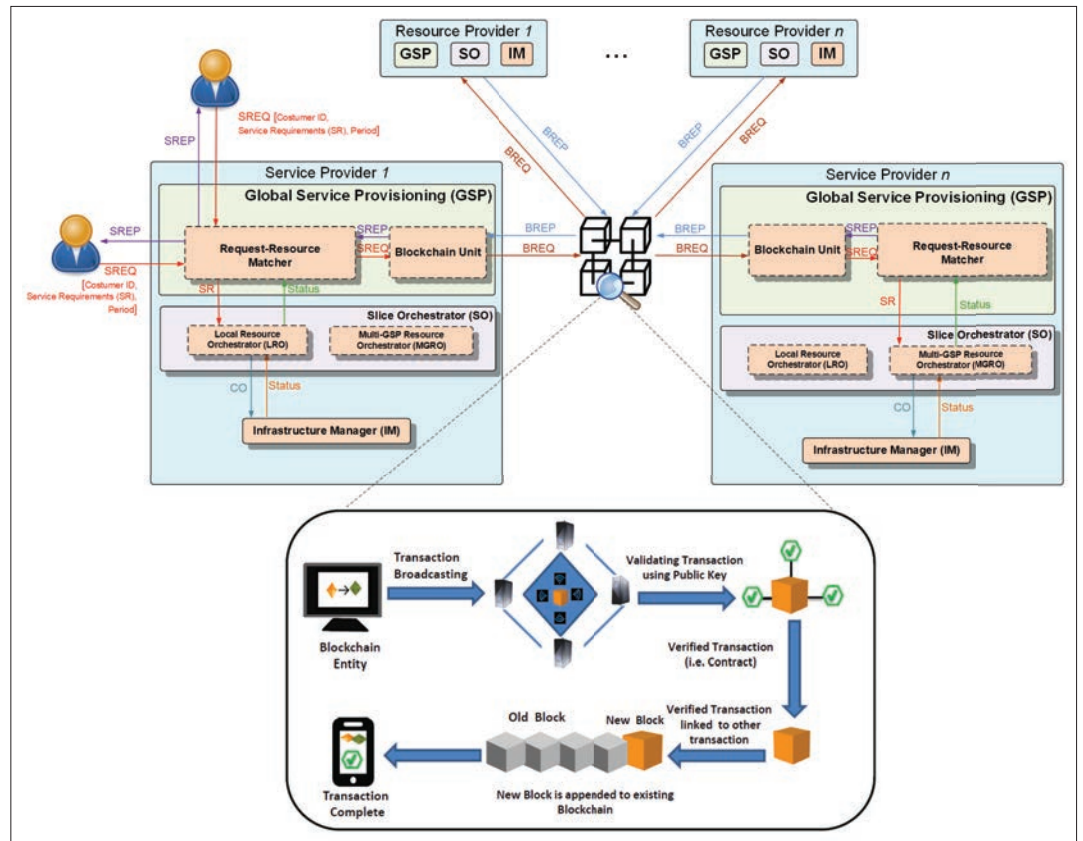


Figure 2. Distributed blockchain network slicing (DBNS) architecture.

inside smart contracts is visible to all nodes. To mitigate these issues, Quorum was proposed. It is an Ethereum-based digital ledger technology (DLT) that provides a layer on top of Ethereum to enable the implementation of private transactions and to enhance the BC network robustness via the deployment of different consensus algorithms. Some of Quorum's most promising features, with respect to our approach, are summarized below:

- *Privacy*: It supports private transactions and private contracts through public/private state separation and uses constellation, a P2P encrypted message exchange, for the transfer of private data among network participants.
- *Peer Permissioning*: It allows node/peer permissioning using smart contracts, ensuring that only known parties can join the network.
- *Higher Performance*: It offers significantly higher performance (100 transactions/s) compared to Ethereum.
- *Alternative Consensus Mechanisms*: It offers multiple consensus mechanisms that are more appropriate for consortium chains.

NETWORK SLICING PLAYERS AND BUSINESS REQUIREMENTS

From the technical point of view, a network slice is an independent logical network that runs on a shared physical infrastructure capable of providing an agreed service quality. From the business perspective, deploying network slicing will enable different players to achieve strategic commercial goals such as increasing return on investmet

(RoI), enhancing network capacity, and expanding network coverage. We identify four players, described as follows:

- *Mobile Network Operator (MNO)*: in charge of deploying and maintaining the physical network.
- *Mobile Virtual Network Operator (MVNO)*: lacks network infrastructure and leases resources from an MNO.
- *Over-The-Top Service Provider (OTT)*: operates on top of a network infrastructure according to a pre-defined set of requirements that are specified in SLAs signed with MNOs/MVNOs.
- *Vertical Industry (VI)*: a set of applications that are specific to an industrial sector (e.g., automotive, fabrication, entertainment). They use the network infrastructure to provide services to end users.

Given that the network infrastructure may belong to either MNOs or MVNOs, addressing the following questions is of paramount importance for these two players to decide whether or not to invest in network slicing: Is it beneficial to their businesses? Can they remain competitive? To this end, the GSMA Network Slicing Taskforce (NEST) specified in [13] three considerations when investing in network slicing. They are summarized here:

- *Complexity*: MNOs/MVNOs need to engage in standardization activities in order to minimize the technical solutions' complexity so that adoption can be made relatively easy.
- *Deployment Scenarios*: All players should agree on the commercial deployment scenarios for network slicing to drive economies of scale.

- **Deployment Cost:** MNOs/MVNOs need to work together to make the cost of deploying network slicing marginal to the broader 5G investment.

While the 3GPP Service Working Group SA1 described in [14] the requirements for various use cases that are related to network slice deployment (e.g., end-to-end asset tracking and wind farm communication networks), the 3GPP Service Working Group SA5 presented in [15] use cases and requirements for management and orchestration of network slicing, including multiple-operator coordination management. Assume that operator A wants to create an end-to-end network slice to support a service across multiple operators (i.e., operators B and C). SA5 enumerates the steps to accomplish this task, summarized as follows:

- Operator A makes the service request to operators B and C asking them to allocate resources (i.e., slices) for the service.
- Operators B and C can either create new slices or use existing ones.
- Operators B and C should provide management data (e.g., performance data) when requested by operator A.
- Operator A is responsible for the management of the end-to-end network slice.

The degree of trust between the different players can have an impact on the entire ecosystem. Indeed, VIs and OTTs must trust MNOs/MVNOs to provide the necessary resources. In addition, MNOs/MVNOs must ensure that the control provided to VIs and OTTs does not allow them to negatively impact their network infrastructure. Finally, MNOs/MVNOs must trust one another to provide the required resources for services spanning multiple administrative domains. For this to happen, additional considerations are needed to provide an appropriate level of control and monitoring.

DISTRIBUTED BLOCKCHAIN-BASED NETWORK SLICING FRAMEWORK

DBNS is a distributed solution that enables network operators to request resource provisioning from other operators in case they cannot fulfill the requested services. Figure 2 depicts its block architecture. It has three components: GSP, slice orchestrator (SO), and infrastructure manager (IM). GSP is in charge of mapping resource provisioning requests to appropriate network slices. This is done through a blockchain-based bidding system that enables all operators having a GSP to participate in the bidding process. It has two modules.

Request-Resource Matcher (RRM): responsible for handling customer requests and mapping them to the suitable network slices. When receiving a customer service request, it first verifies whether the current provider can fulfill the service by probing the SO. If not, the RRM plays the role of a broker by broadcasting the service provisioning request over the blockchain and selecting the best bid among the ones received from the other operators.

Blockchain Unit (BU): responsible for writing resource provisioning requests to the blockchain and fetching their corresponding bids. The resource provisioning requests, called bid

requests, are advertised in plain text, while the matching bids, called bid replies, are encrypted using the public key of the requesting operator.

The SO provides end-to-end network slice orchestration either locally or across multiple administrative domains. It has two modules.

Local Resource Orchestrator (LRO): provides an abstract view of the underneath infrastructure to GSP with the help of IM. It makes decisions about creating new slices and updating existing ones to include virtual/physical network functions. It also keeps track of utilization status of the various slices and monitors the QoS metrics to comply with SLA requirements.

Multi-GSP Resource Orchestrator (MGRO): allocates the appropriate resources (i.e., by creating new slicing or upgrading existing ones) for the requested service if the bid reply is selected and monitors the QoS metrics to ensure that SLA requirements are properly met. It also enables seamless service transition from one operator to another and provides usage reports to be used for billing purposes.

Finally, IM manages the underlying physical/virtual infrastructure. It provides support for slicing, enforces slice requirements from the SO, and provides infrastructure monitoring and analysis services.

To allow for a secure and efficient service provisioning, DBNS has five phases.

Phase 0 – Setup: An operator willing to join the DBNS framework is required to obtain a public key certificate from a trusted third party (TTP) such as a commercial certification authority (CA). This certificate is then shared with existing partners (i.e., operators that are already members of the DBNS framework) to ask to join the framework. Upon agreement, partners virtually approve the operator's request. This latter then signs the *GSP-based onboarding* digital contract, which specifies the activities that can be performed and the rules to abide by. This includes requesting resource provisioning from other partners by broadcasting bid requests and sending bid replies. If the placed bid is selected, partners have the obligation of fulfilling the requested service according to the agreed SLA requirements.

Phase 1 – Consumer Request: To avail of the various services, consumers send service requests (SREQs) to their network operators. Each SREQ contains the consumer ID, the service requirements (e.g., bit rate, loss range, latency), a time period during which the service is needed, and the location where the service is to be provided. When received, the RRM forwards the service requirements to the SO. This latter probes the LRO to verify whether the service can be fulfilled locally. With the help of IM, the LRO sends a *status* message to the RRM. If the service can be supported, *Phase 3* is executed. Otherwise, *Phase 2* is carried out.

Phase 2 – BC-Based Bidding: If the service cannot be fulfilled locally, the RRM encapsulates an SREQ into a bid request (BREQ) containing the request ID along with the ID of the requesting operator and forwards it to the BU. This latter writes BREQ into the blockchain. DBNS makes use of a *permissioned* blockchain to solicit bids from different partners in the ecosystem.

The degree of trust between the different players can have an impact on the entire ecosystem. Indeed, VIs and OTTs must trust MNOs/MVNOs to provide the necessary resources. In addition, MNOs/MVNOs must ensure that the control provided to VI and OTTs does not allow them to negatively impact their network infrastructure.

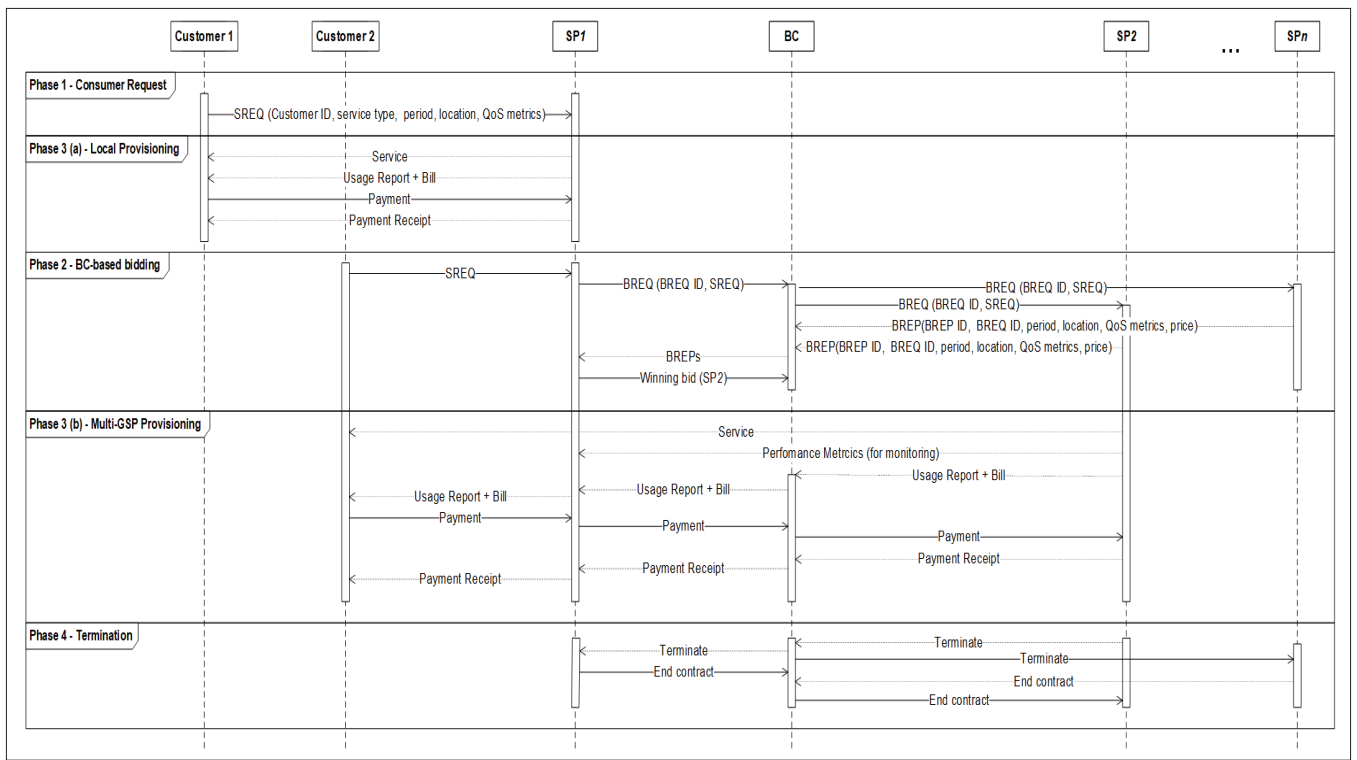


Figure 3. DBNS sequence diagram illustrating the various phases.

Parameters		Value
Simulator		NS-3 version 3.26
Duration		300 seconds
Number of UEs		10 UEs
BS_A & BS_B	LTE eNodeB Antenna Model	Isotropic Antenna Model
	LTE Path Loss Model	Friis Propagation Loss Model
AP_C	Wi-Fi Standard	802.11ac
	Wi-Fi Setting	VHT Mode (80MHz, MCS 9)
	Wi-Fi Path Loss Model	Log-dist. Propagation Model

Table 1. Simulation setup parameters.

Each partner’s BU fetches the BREQ from the blockchain and hands it to the RRM, which gets the encapsulated SREQ. The RRM then sends the service requirements to the MGRO to verify whether the service can be supported. If yes, the RRM makes an offer (i.e., includes the minimum and maximum supported service requirements along with the price), encapsulates it into a bid reply (BREP), signs it with the public key of the requesting operator, and forwards it to the BU. The signature is used to prevent other providers from obtaining any details about the bid, ensuring a fair and transparent competition. The BU adds the BREP transaction to the blockchain. The BU of the requesting operator fetches all BREPs and forwards them to the RRM. This latter decrypts them using its private key, gets the different offers, and chooses the one that perfectly suits the needs of its customer. At this stage, the requesting operator executes a “smart contract” on the BC, which locks in the customer and the operator providing the service (i.e., the bid winner). The smart contract is used to hold both parties accountable in terms of their respective obligations.

Phase 3 – Service Provisioning: In the case of local service provisioning (Phase 3a in Fig. 3),

the RRM asks the LRO to allocate the necessary resources to meet the service requirements and to start the service. Once the service is completed, the RRM generates the usage report along with the bill and sends it to the customer. If the service is to be provided by another partner, the corresponding RRM asks the MGRO to allocate the necessary resources and to start the service. Once the service is completed, the RRM generates the usage report along with the bill and adds them to the blockchain (i.e., encrypted with the public key of the requesting operator). The BU of the requesting provider fetches the report along with the bill and forwards them to the RRM, which transmits them to the consumer (Phase 3b in Fig. 3).

Phase 4 – Termination: An operator can leave the DBNS framework voluntarily by broadcasting a terminate message via the blockchain. Before departing, the operator should make sure that no new service provisioning requests are accepted and that all ongoing services are completed. Upon approval, partners then terminate the operator’s GSP-onboarding contract and write the transaction on the blockchain (Phase 4 in Fig. 3). Partners can also impose sanctions on an operator in the case of recurrent noncompliance with SLA requirements. In this regard, bid request requirements are compared against usage summary reports, and penalties (e.g., boycotting the bid replies of the inspected operator for a period of time) may apply for reported transgressions.

CASE STUDY: ENTERTAINMENT SERVICES

This section presents a simulation-based performance evaluation of DBNS. The case study considered replicates a real-life situation where a customer U_A of a network operator SP_A streams Youtube videos with resolutions 720p and 1080p,




Figure 4. Simulation results of the various requested services with and without use of DBNS: a) average throughput; b) average delay and average loss rate.

plays high resolution 3D games (i.e., with a 2K video component), and watches 4K Netflix movies. These entertainment services have high bit rate requirements of 2 Mb/s, 4 Mb/s, 15 Mb/s, and 25 Mb/s. U_A is connected via base station BS_A .

We assume that BS_A has a dedicated network slice for entertainment services that is already overloaded as it is serving multiple concurrent users. In the same area, operators SP_B and SP_C offer services via base station BS_B and access point AP_C , respectively. Two scenarios are considered:

- SP_A only uses its encumbered network slice (between BS_A and U_A) to fulfill U_A 's service requests. This scenario is labeled *Without DBNS*. Given that existing designs either have no performance evaluation or are not suitable for comparison with our proposed architecture due to the use of auxiliary technologies, the *Without DBNS* approach is chosen as a benchmark.
- SP_A employs the proposed DBNS approach to request service provisioning from SP_B and SP_C . This scenario is labeled *With DBNS*. SP_A selects SP_C 's bid as it has enough network resources to be allocated to U_A . The service is fulfilled via the link between AP_C and U_A .

Other simulation parameters are summarized in Table 1.

Figure 4 shows how, by employing the proposed DBNS, the supported throughput is equivalent to the bit rate required by the various services (i.e., 2 Mb/s for 720p, 4 Mb/s for 1080p, 15 Mb/s for 2K, and 25 Mb/s for 4K). Figure 4a shows that without deploying DBNS, only the bit rates of 720p and 1080p services can be supported due to the lack of sufficient resources at BS_A . Figure 4b depicts the average delay and loss rate when using both approaches. Using DBNS, the loss rate is almost 0, while the average delay ranges between 20 and 25 ms for the different services. Without DBNS, the average delay is between 19 and 22 ms, which matches the one incurred by DBNS, particularly for 2K and 4K services. This is because the average delay in this case only considers the supported throughput for both services (i.e., 8 Mb/s each) and not the

required one. Furthermore, the loss rate is exorbitant for 2K (i.e., 50 percent) and 4K (i.e., 70 percent) services, which significantly affects the users' quality of experience.

CONCLUSIONS

This article discusses network slicing technology, describing its major players and business requirements in the context of the 5G technology. It then introduces the DBNS framework, which promotes dynamic resource leasing between different providers to support high performance for end-to-end services. It has two key components and their associated mechanisms: GSP, which provides admission control for incoming requests by means of a blockchain-based bidding system, circumventing the lengthy process of setting up MoUs; and the MGRO, which supports seamless service transition from one entity to another along with monitoring QoS metrics to ensure that SLA requirements are met.

We believe that DBNS is a fundamental leap forward as it enables network operators to move toward a more flexible and on-demand way of leasing and procuring network infrastructure and assets to fulfill services in situations where infrastructure cannot meet demand, and to provide more dynamic capacity when demand spikes are experienced (i.e., "pop-up" networks for large sporting events). This can help network operators decrease both their CAPEX and OPEX costs and increase their revenues.

Future research directions could consider the analysis of other context-aware, multi-tenant use cases and design of innovative resource allocation techniques considering the isolation requirement. In addition, the evaluation of DBNS from business, policy, and legal perspectives could be realized. Finally, blockchain performance optimization in terms of transaction rate and scalability can be evaluated through the use of different data structures as well as deployment of multiple chains or channels.

ACKNOWLEDGMENT

The Science Foundation Ireland (SFI) support, co-funded by the European Regional Development Fund, via research grants 16/SP/3804

(Enable) and 12/RC/2289_P2 (Insight) is gratefully acknowledged.

REFERENCES

- [1] 3GPP, "Feasibility Study on New Services and Markets Technology Enablers for Critical Communications," Tech. Rep. 22.862, R. 14, 2016.
- [2] K. Samdanis, X. Costa-Perez, and V. Sciancalepore, "From Network Sharing to Multi-Tenancy: The 5G Network Slice Broker," *IEEE Commun. Mag.*, vol. 54, no. 7, July 2016, pp. 32–39.
- [3] V. Sciancalepore et al., "Mobile Traffic Forecasting for Maximizing 5G Network Slicing Resource Utilization," *IEEE INFOCOM*, May 2017, pp. 1–9.
- [4] V. Sciancalepore, F. Cirillo, and X. Costa-Perez, "Slice as a Service (SaaS) Optimal IoT Slice Resources Orchestration," *IEEE GLOBECOM*, Dec 2017, pp. 1–7.
- [5] T. Taleb et al., "On Multi-Domain Network Slicing Orchestration Architecture and Federated Resource Control," *IEEE Network*, vol. 33, no. 5, Sept./Oct. 2019, pp. 242–52.
- [6] G. Kibalya et al., "A Reinforcement Learning Based Approach for 5G Network Slicing Across Multiple Domains," *Proc. 15th Int'l. Conf. Network and Service Management*, Oct. 2019, pp. 1–5.
- [7] A. Devlic et al., "NESMO: Network Slicing Management and Orchestration Framework," *Proc. IEEE ICC Wksp.*, May 2017, pp. 1202–08.
- [8] M. A. Togou et al., "A Hierarchical Distributed Control Plane for Path Computation Scalability in Large Scale Software-Defined Networks," *IEEE Trans. Network and Service Management*, vol. 16, no. 3, Sept. 2019, pp. 1019–31.
- [9] H. Yang et al., "Blockchain-Based Hierarchical Trust Networking for Jointcloud," *IEEE Internet of Things J.*, vol. 7, no. 3, 2020, pp. 1667–77.
- [10] H. Yang et al., "Distributed Blockchain-Based Trusted Multi-Domain Collaboration for Mobile Edge Computing in 5g and Beyond," *IEEE Trans. Industrial Informatics*, 2020, pp. 1–1.
- [11] H. Tewari, "Blockchain Research Beyond Cryptocurrencies," *IEEE Commun. Standards Mag.*, vol. 3, no. 4, Dec. 2019.
- [12] M. Belotti et al., "A Vademecum on Blockchain Technologies: When, Which, and How," *IEEE Commun. Surveys & Tutorials*, vol. 21, no. 4, 2019, pp. 3796–3838.
- [13] GSMA, "Network Slicing Use Case Requirements," Apr. 2018; <https://www.gsma.com/futurenetworks/wp-content/uploads/2018/07/Network-Slicing-Use-Case-Requirements-fixed.pdf>.
- [14] 3GPP, "Feasibility Study on Business Role Models for Network Slicing," Tech. Rep. 22.830, Rel. 16, Dec. 2018.
- [15] 3GPP, "Study on Management and Orchestration of Network Slicing for Next Generation Network," Tech. Rep. 28.801, Rel. 16, Jan. 2018.

BIOGRAPHIES

MOHAMMED AMINE TOGOU is a postdoctoral researcher with the Performance Engineering Laboratory at Dublin City University (DCU), Ireland. He received B.S. and M.S. degrees in computer science and computer networks from Al Akhawayn University, Ifrane, Morocco, and a Ph.D. degree in computer science from the University of Montreal, Canada. His current research interests include 5G networks, SDN-NFV, network slicing, and IoT.

TING BI is a postdoctoral researcher with the Performance Engineering Laboratory at DCU. He received a B.Eng. in software engineering from Wuhan University, China, in 2010, and M.Eng. and Ph.D. degrees in telecommunications from DCU in 2011 and 2017, respectively. His research interests include mobile and wireless communications, multimedia, and multi-sensory media streaming.

KAPAL DEV is a research fellow with the CONNECT Centre, School of Computer Science and Statistics, Trinity College Dublin. He was awarded a Ph.D. degree by Politecnico di Milano, Italy, in July 2019. His research interests include blockchain, 5G beyond networks, and artificial intelligence. Previously, he worked as a 5G junior consultant and engineer at Altran Italia S.p.A, Milan. He is coordinating two Erasmus+ International Mobility projects.

KEVIN MCDONNELL is an architect at the Huawei Ireland Research Centre. His recent work focuses on autonomous network operations using knowledge management, intent, and machine learning. He serves as workstream leader for Autonomous Networks in the TM Forum, and participates in standardization activities in 3GPP and GSMA. He is an active contributor to the ONAP project and is company leader for EU Horizon 2020 Project 5GTANGO.

ALEKSANDAR MILENOVIC is a Huawei chief architect for Autonomous Network Operations and director of the Communications Technology Lab in the Huawei Ireland Research Center. He has nearly 20 years of experience in developing products in the area of telco operations and maintenance (O&M). He is leading research in intelligent automation of telco operations.

HITESH TEWARI is an assistant professor in the School of Computer Science and Statistics at Trinity College Dublin. His research interests lie in the areas of network security and applied cryptography. In recent years he has been actively working in the area of distributed ledger technologies.

GABRIEL-MIRO MUNTEAN (gabriel.muntean@dcu.ie) is an associate professor with the School of Electronic Engineering, DCU, and co-director of the DCU Performance Engineering Laboratory. He was awarded a Ph.D. degree by DCU in 2004. His research interests include quality-, energy-, and performance-related issues of rich media content delivery. He is an Associate Editor of *IEEE Transactions on Broadcasting* and Multimedia Communications Area Editor of *IEEE Communication Surveys & Tutorials*. He coordinated the EU Horizon 2020 project NEWTON.