Blockchain-driven mobile data access towards fully decentralized mobile video trading in 5G networks

Dionysis Xenakis
Fogus Innovations and Services P.C.
Email: dionysis@fogus.gr

Ioannis Zarifis, Pantelis Petrogiannakis
Nessos Technologies S.A.
Email: {izarifis, ppetrogian}@nessos.gr

Anastasia Tsiota, Nikos Passas
University of Athens
Email: {atsiota, passas}@di.uoa.gr

***Abstract*—The today's heterogeneous wireless network infrastructure is managed by different stakeholders that allow mobile data access to a prescribed set of end users. This network access model not only enforces end users to accept network-driven service provisioning but also raises critical burdens towards flexible utilization of the numerous 5G network assets (content, storage, radio resources, etc.). In this paper, we investigate on-the-fly user-driven network-assisted mobile video content delivery in 5G and Beyond networks using blockchain technologies enabling multi-million transactions per second. To this end, we present the RE-CENT crypto-currency platform enabling infrastructure providers, end users and content providers to be actively engaged in the blockchain consensus. Detailed discussions on the implementation of the proposed platform through smart contracts are provided and analytical results are used to quantify performance gains of the proposed mobile data access model.**

***Index Terms*—5G networks, mobile video content, blockchain, network asset trading, crypto-currency platform.**

I. Introduction

The 5th Generation mobile data network encompasses exciting new radio capabilities enabling high-data rates with small over-the-air delay [1]. Nonetheless, the availability of such high-data rate links is not guaranteed in high peak periods due to the random topology characterizing 5G networks (e.g. user-installed base stations), the diverse radio access technologies coexisting under the same network architecture (e.g. mmWave, cellular, WiFi) and the obsolete service models governing access to the heterogeneous wireless network infrastructure.

Although notable steps have been recently made towards over-the-top (OTT) service provisioning and infrastructure sharing across different mobile network operators (MNOs) [2], mobile data access of end users is still governed by fixed-term contracts with a given MNO, or short-term access coupons granted through eponymous payments using fiat money, or short-term "free-of-charge" access dictating user consent on the collection and processing of their personal data (e.g. enforcing log in with personal email or social media accounts). Existing models regulating end user access in pre-5G systems are evidently inflexible, creating a critical burden towards seamless and fully-personalized content consumption.

In parallel, triggered by the successful Bitcoin platform [3], blockchain-based system design has opened up new horizons towards the enforcement of provably-secure distributed consensus across non-trusted peers. The consensus protocol is the core component of a crypto-currency system, ensuring that all participating nodes agree on a common transaction history that is serialized and crystallized due time into consecutive blocks that form the distributed ledger, a.k.a. the blockchain. Early consensus protocols are based on a the participation of the consensus nodes in a solution searching process, known as Proof of Work (PoW), leading to high energy consumption and very low transaction (Tx) throughput [4].

To overcome this, new consensus mechanisms have emerged [5], based on Proof-of-Useful-Work (e.g. Primecoin), Proof-of-Space (e.g. SpaceMint), Proof-of-Storage (e.g. Filecoin), Proof-of-Elapsed-Time (e.g. Hyperledger), Proof-of-Stake (e.g Ouroboros), Proof-of-Authority as well as hybrid PoW and PoS. PoS and PoA in particular, have attracted surge of interest lately, due to their capability to provide energy-saving alternatives to PoW with relatively low-cost. PoS consensus selects block sealers (similar to Bitcoin miners) based on their stakes (balance) in the system, whereas PoA consensus selects authorizes a very small set of nodes to seal blocks, a.k.a. validators, that either provide proofs of their identity publicly (e.g. ID, residence address), or acquire a special license by authorized notaries [6]. PoA has recently gained momentum due to the release of the Parity Aura consensus protocols in Ethereum [7], and the numerous PoA-based services made available by key industry players, e.g. Microsoft and Amazon.

In this paper, we propose the use of crypto-currency platforms that are specifically designed to turn the today's evidently under-organized and vastly heterogeneous mobile data network, where different operators, regional network /content providers, user-installed access points and end terminals, share no interest in improving the networking experience of end users belonging outside their subscriber's whitelist, to a fully decentralized, dynamic and competitive (by consensus) market where different stakeholders have clear incentives to improve content consumption of mobile users within coverage.

Section II details a radically new model for flexible mobile data access, enabling infrastructure owners to trade network assets in a fully decentralized, reliable and secure fashion. Section III draws implementation details of a specialized crypto-currency platform to this end, specifying system architecture and detailing implementation through the deployment of smart contracts that allow i) 5G-specific network consensus and ii) high transaction throughput using off-chain payments. Section IV quantifies performance gains of the proposed mobile data access model while Section V concludes the work.

II. The RE-CENT Mobile Data Access Model

End users interested in consuming video content from OTT service providers located in the far Internet, should go through the mobile data network infrastructure having as a entry key pre-cached access credentials (subscriber ID, IMSI, social media account, passwords, etc.) granting them access to a limited set of network islands. Focusing on the exchange of mobile video content, which will account for over 74% of the total mobile data traffic by 2024 [8], in the sequel we propose and investigate the potential of a new network access model that we term as *REsource sharing model for user-CENTric mobile video content delivery (RE-CENT)*. The RE-CENT model assumes a flat resource trading architecture where network nodes act both as network asset agents (servers) and consumers (clients). Every RE-CENT node is holder of at least one public address (wallet ID) that enables on-the-fly consumption or delivery of mobile video content.

A. Network asset advertisement and discovery

This phase can be implemented by using wireless network-level protocols (e.g. physical network and OTT service discovery), or by using on-chain methods (e.g. every agent advertises available resources and costs using on-chain transactions [9]). Network-level asset advertisement and discovery enables clients to communicate their mobile video content request locally (e.g. using a URL) and receive targeted service offers. In this manner servers can adapt their pricing policy based on their current connectivity and availability of local content (e.g. lower fees for cached content). In the contrary, on-chain asset advertisement enables transparency and immutability at the cost of a large volume of on-chain transactions necessary to keep service offers up to date (on a per server basis). Filtering of relevant asset offerings on the global (public) ledger is also challenging due to user mobility. Given the extremely high number of potential service peers and the number of video files available globally, both of which are in the order of billions [8], the RE-CENT crypto-currency platform assumes the deployment of network-level asset advertisement and discovery.

B. Service negotiation, parameterization, pricing, charging

Having discovered available network assets, RE-CENT clients and servers subsequently negotiate service parameters.

1) How the client communicates to the server the content request: The RE-CENT service request may include a specific URL; however, more sophisticated techniques can also be deployed (use of specific tags, or video authors, etc.).

2) How will the client estimate its Quality of Service (QoS) / Quality of Experience (QoE) requirements and how will the server estimate the minimum QoS/QoE it can guarantee: The RE-CENT client should specify the value of key performance indicators (KPIs) affecting the QoS/QoE of the envisaged service. The RE-CENT server will adapt service offers and costs according to the QoS/QoE targets set by the RE-CENT client. Service peers can encompass existing state-of-the-art mechanisms for accurate QoS/QoE estimation to this end [10].

3) How will the server conclude on the price of a custom service: RE-CENT servers shall deploy their own pricing strategies in a competitive and open network asset market, taking into account: a) the components defining the real cost of the end-to-end (e2e) content delivery chain (local communication and storage, backhaul and core network support), b) the pricing policy of nearby competitors and c) the client's ambition to increase its market share (e.g. asking low fees for new users). The cost for delivering locally available (cached) content can be significantly lower as compared to fetching the requested content from the far Internet, linking the effectiveness of content caching to the pricing logic [11].

4) How will the platform guarantee reliable content consumption and secure payments: The hash of all parameters defining the service instance can be included in the transaction payload to identify the service and formalize the agreement. However, advanced mechanisms should be considered to protect user privacy in this case. Service setup can be formalized using smart contracts and on-chain transactions. However, such an approach would further increase the already large volume of transactions without necessarily enforcing sufficient trust across the service peers. Alternatively, the RE-CENT service peers should agree in advance on a specific payment plan, both in terms of timing (e.g. based on video time, or per MB basis) and amount of intermediate (micro-)payments. The timing and amount of intermediate payments can be adjusted according to the level of trust among the two peers.

C. Online service management

RE-CENT service management should be fully aligned with heterogeneous network access, making the user fully responsible for predicting potential service discontinuation. For example, early prediction of user mobility could trigger the establishment of multiple service instances which, combined with user-controlled network-assisted mobility management [12], can mitigate unnecessary service interruption. Adaptive video streaming can also enable end terminals to control e2e video delivery in response to service quality fluctuations [13].

III. The RE-CENT Platform

A. System Architecture

We consider a multi-tier HWN, where each network tier is composed by wireless networking elements (WNEs) that support the same technology and operate in the same frequency band. On-top of the HWN infrastructure, we consider a flat system architecture where WNEs form a flat logical architecture for network asset trading: the RE-CENT service domain (Fig. 1). Every RE-CENT service node uses at least one public address to deliver and consume mobile video content by trading: i) local content, ii) end-to-end connectivity to the Internet, and iii) (relay) connectivity to other nodes.

RE-CENT nodes form the RE-CENT consensus network to disseminate transactions and maintain the RE-CENT blockchain (Fig. 1). Using the RE-CENT delegated PoS (DPoS) mechanism (section III.C), the consensus network authorizes a limited number of RE-CENT nodes (termed as

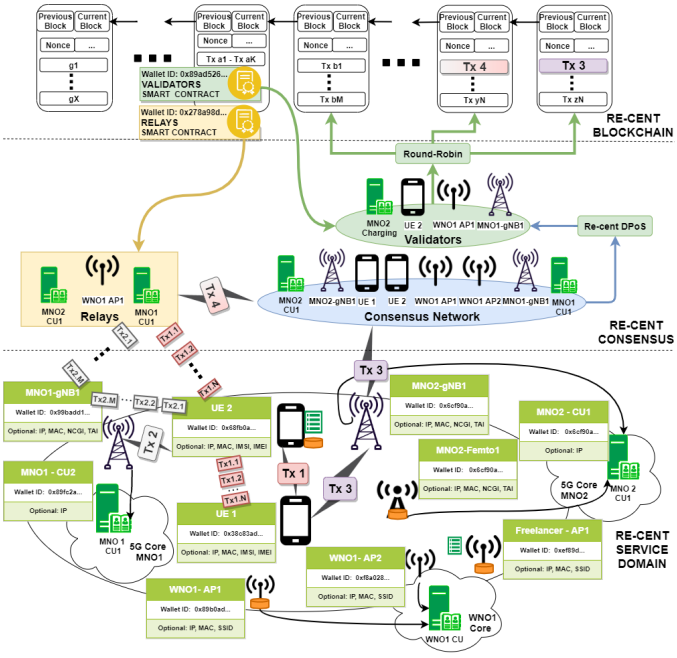

Fig. 1. RE-CENT System Architecture

validators) to seal new blocks taking series in a round-robin fashion. Validators are elected for only a prescribed time period that is termed as *epoch* (measured in blocks). All parameters and rules affecting the validation process, are publicly known and implemented by the *Validators* SC (VSC) that is deployed in the early blocks of the RE-CENT blockchain (section III.C). The use of DPoS consensus along with round-robin block sealing from a very small group of validators aims to the attaining a very high transaction throughput.

Smart contracts (SC) are self-executing scripts that reside on the blockchain and enable general-purpose computations to occur on chain. Every SC possesses a unique address and transactions containing data can be sent to that address to trigger the execution of the SC code. Ethereum is the mains platforms enabling the execution of SCs [14]. In the sequel, we assume the Ethereum Aura platform for system prototyping purposes [7] (PoS consensus is not yet available by Ethereum).

Taking into consideration the enormous number of service peers that could potentially participate in this process, support of the RE-CENT access model requires multi-million transaction throughput per second. Since the transactional capacity of existing cryptocurrency platforms is in the order of hundreds transactions per second [5], novel consensus protocols and system design approaches are required to this end. To this end, founded on the concept of payment channels [15], the RE-CENT service architecture enables RE-CENT nodes to use payment *relays* in order to attain low-cost but nearly instant *off-chain transactions*. Relays are considered as trusted payment intermediaries between the service peers and the RE-CENT blockchain, with their trust enforced by the RE-CENT *Relays* SC (RSC - Section III.D). Off-chain transactions are legitimate transactions that are crypto-graphically signed by the RE-CENT client that are buffered in the relays and are released using specific mechanisms enforced by the RE-CENT client and the RSC logic. Relays can be used to aggregate intermediate (micro-)payments between service peers, enabling the system's transaction throughput to scale with the number of service peers. Depending on their effectiveness, relays may further aggregate payments between multiple service peers only submitting the final outcome (balance) of the participants in the blockchain (allowing infinite transaction throughput).

B. Interaction across peers and service flow

1) Service flow using direct payments: Having concluded on the RE-CENT service setup, the RE-CENT service peers shall initiate the video delivery service immediately. Following the agreed payment timeplan, the RE-CENT server will issue payment requests that the RE-CENT client should sign and send back. In this manner, the RE-CENT client is not required to have other connections to the Internet while the RE-CENT server is fully responsible for submitting signed transactions to the consensus network. Fig. 1 illustrates how a direct payment Tx3 is included in the RE-CENT blockchain.

2) Service flow using relay payments: In this scenario, the RE-CENT service peers shall agree on the relay, taking into account the maximum delay guarantee offered by the relay and the maximum delay tolerated by the RE-CENT client (prior to the on-chain update of its balance by the relay). When the RE-CENT server is about to issue a payment request to the RE-CENT client, it asks from the relay to confirm the hash of the last block finalized in the RE-CENT blockchain (using off-chain mechanisms) and sign a transaction including the maximum tolerable delay by the RE-CENT client. Assuming that the relay is enforced to provide the correct answer by the RSC, the RE-CENT server issues a payment request to the RE-CENT client including in the payload signed by the relay. To avoid service termination, the RE-CENT client will sign the transaction and send it back to the RE-CENT server that will then forward the signed transaction(s) to the relay. The relay shall validate that the balance of the RE-CENT client is sufficient and shall notify the RE-CENT server accordingly.

The RE-CENT server shall communicate with the client over-the-air (in parallel with the video service) and shall reach the relay using web services. The maximum tolerable delay specified by the RE-CENT client can be used to schedule off-chain payments and enforce RSC-driven relay compliance with the RE-CENT server requirements (section III.D).

C. Validators Smart Contract Overview

1) VSC main parameters: The *Validators SC* (VSC) deploys the governance model of the RE-CENT blockchain system. Assuming that the current epoch is denoted by e, the key parameters of the VSC are: i) the minimum, maximum and current number of validators, denoted by V_{max}, V_{min} and $V[e]$, respectively ii) the minimum, maximum and current epoch duration, denoted by B_{max}, B_{min}, $B[e]$, respectively, iii) the number of blocks before the end of an epoch until which the staking process is concluded, denoted by T_V, iv) the

emission rate R of new coins generated per block, v) a table of disinflation rates d_t that are applied to the emission policy in time intervals $t \in \mathcal{D}$ (measured in blocks), vi) the minimum transaction fee for direct payments, denoted by c_{min}, vii) the current penalty fund for validators and the penalty percentage applied to the validator witness balance, denoted by $P_V[e]$ and $p_W[e]$, respectively, viii) the minimum stake required for validator candidates, denoted by $M_V[e]$, ix) the minimum stake required for validator witnesses, denoted by $M_W[e]$, x) the free service staking cost per MB, denoted by $f_V[e]$ and xi) the number of consecutive epochs required to amend some VSC parameter, denoted by $C[e]$.

The minimum and maximum values of V_{max}, V_{min}, B_{max}, B_{min} as well as the values of e, d_t, $\mathcal{D}$, c_{min} and T_V are hard-coded in the configuration (.config) file of the RE-CENT blockchain. All the remainder VSC parameters can be configured if: i) at least $2/3$ of the validators in a consecutive number of at least $C[e]$ epochs has voted in favor of the amendment and ii) the revision of the value is in the range of a predefined step that is hard-coded in the blockchain .config file (e.g. +/-1 for $V[e]$ and $B[e]$). The aforementioned amendment mechanism enables fine-tuning of the VSC parameters in line with the status of the RE-CENT consensus network.

The penalty fund of validators shall be a function of the amount of newly minted coins per epoch. The penalty fund of validator witnesses (v-witnesses) shall be a percentage of their total balances (e.g. $p_W = 3\%$ per epoch). $M_V[e]$ and $M_W[e]$ should be adjusted carefully to avoid an infinite number of candidate validators and validator witnesses. If the maximum number of candidates is reached in the VSC regitry (due to size limitations), a new candidate can replace an existing one by submitting signed transactions with higher v-witness stakes.

2) The RE-CENT DPoS mechanism: The staking - validator election process for epoch e starts with epoch $e-1$ and concludes T_V blocks before epoch e. In the sequel we term epoch $e-1$ as the *staking epoch* and epoch e as the *target epoch*. A RE-CENT node is considered as candidate validator if it sends to the VSC: i) signed transaction locking the penalty fund $P_V[e]$ to the VSC, ii) a reward that will be shared among the v-witnesses according to their share during the target epoch, referred to as the *v-witness reward balance* of the validator, and iii) a fixed transaction fee that the validator will charge per on-chain transaction, denoted by $f_V[e](> c_{min})$.

RE-CENT nodes are enabled to provide free-of-charge service to v-witnesses that stake in favor of a given candidate validator (in case it gets elected). The so-called free-of-charge RE-CENT servers are required to time-lock to the VSC a fee of $f_V[e] \cdot X_V[e]$ RE-CENT coins, where $X_V[e]$ is the value of free MBs offered by the RE-CENT servers per v-witness. This fee serves as a penalty fund in case the free-of-charge RE-CENT server fails (or refuses) to provide the promised service to v-witnesses. Free-of-charge service staking aims to attract more v-witnesses and incentivizing active RE-CENT users to participate in the DPoS election mechanism.

RE-CENT nodes that are interested in acting as v-witnesses, retrieve the list of candidate validators and free-of-charge servers by the VSC. Accordingly, they time-lock a fixed percentage p_W of their current balance to the VSC, indicating the preferred candidate validator(s). The accumulated v-witness stakes for a given candidate validator are termed as *coin staking balance* of the validator. RE-CENT validators prioritize processing of transactions including addresses of their v-witnesses, to give additional incentives for RE-CENT nodes to act as v-witnesses. Validators' penalty funds, v-witness rewards, v-witness stakes and free-of-charge penalty fees are time-locked for both the staking and the target epoch.

3) The RE-CENT validation process: T_V blocks before the start of epoch e, the list of candidate validators is ordered based on the sum of the penalty funds and fees committed per validator and the first $V[e]$ candidates are elected as validators. A limited number of elected validators (e.g. $V[e] = 20$) has been shown to be a highly efficient low-cost alternative to PoW attaining a high degree of decentralization [16].

V-witness and free-of-charge RE-CENT server operation. After the start of epoch e, v-witnesses of elected validators can withdraw their share from the respective v-witness reward balance and utilize free-of-charge services. RE-CENT servers providing free-of-charge service will issue empty payment requests to v-witnesses, using signed responses as proofs of their service to the VSC (if necessary). If a v-witness submits to the VSC immutable proofs of a free-of-charge service request that has not been implemented (e.g. association messages accompanied by the hash of the last $2V[e]$ blocks), then the VSC will burn part of the free-of-charge penalty fee submitted by the respective free-of-charge server (i.e. proportional to the service request that can be proved).

Actions related to the validator operation By inserting legitimate transactions to a new block, validators receive relevant transaction fees $f_V[e]$ and newly minted coins $R \cdot d_t$. If a validator fails to seal a new block during its turn (e.g. due to persistent Internet failures), the VSC will burn a portion of the validator's penalty fund and the v-witness coin staking balance. This portion shall increase rapidly with the number of block sealing failures according to $2^{l_v - 1} \cdot \frac{P_V B[e]}{V[e]}$. By the end of the target epoch, validators, v-witnesses and free-of-charge servers associated with non-empty balances can withdraw the remaining penalty fund from the VSC. However, if the coin staking balance of a given validator gets empty before the end of the epoch (due to the penalties received), a *validator replacement* process is triggered. The highest-ranking non-active validator candidate belonging in the ordered list of non-elected validators will replace the banned validator (recall that all stakes and fees are time-locked for both the staking and target epochs). Until the new validator starts to seal blocks, the previous in turn validator shall issue empty blocks in the turn of the banned validator. Validator replacement also takes place if a validator seals erroneous transactions, or blocks outside its ordered sequence. In such occasions, all penalty funds attached to the validator in the VSC are burned at once and the validator replacement process is triggered immediately.

D. Relay Smart Contract Overview

The RSC defines a mechanism to authorize RE-CENT nodes to act as relays, termed as *relay licensing mechanism*, and specifies necessary reward/penalty mechanisms to enforce relay compliance of relays with their license parameters. Relays incur low transaction fees to RE-CENT nodes as compared to direct payments, by buffering RE-CENT transactions and reducing the number of (expensive) on-chain payments. Use of relays is optional; however, high transaction fees for direct payments (¿c_{min}) shall indirectly enforce its use by RE-CENT nodes. Relays are assumed to deploy their own off-chain mechanisms to collect off-chain payments, e.g. web services.

1) RSC main parameters: The RSC includes the following key parameters: i) *relay user registry*, recording the amount of coins attached to a given relay, ii) *relay registry*, recording all parameters necessary for monitoring and validating the operation of relays, iii) *staking registry*, recording parameters necessary for licensing and penalties, iv) four *tariff tables* specifying relay license costs based on the i) maximum number of attached RE-CENT nodes (relay users), ii) maximum amount of attached coins and iii) mean transaction throughput.

Additional parameters of the RSC include: i) number of blocks within which relay licensing should conclude, denoted by T_R, ii) duration of the block period within which the RSC regulates the transaction throughput of relays, denoted by k_R, iii) mean transaction throughput for RSC transactions on the RE-CENT blockchain, denoted by $\bar{R}$, iv) penalty for delayed payments, denoted by v_R, v) percentage threshold of the relay penalty fund below which the relay's license is revoked, denoted by x_R, and vi) free-of-charge service staking cost per MB, denoted by $f_R[e]$. Estimating $\bar{R}$ is necessary since the block size varies during an epoch (adjusted by validators). T_R and k_R are specified in the RE-CENT .config file.

For a tagged relay r, the *relay registry* shall record: i) transaction fee per off-chain transaction, denoted by o_r, ii) current and maximum number of attached nodes, denoted by N_r and $N_{r,max}$, respectively, iii) current and maximum amount of attached coins, denoted by U_r and $U_{r,max}$, respectively, iv) mean (allowed) (on-chain) transaction throughput, denoted by $\bar{R_r}$, iv) counter measuring on-chain transactions for the last period of k_R blocks, denoted by txc_r, v) index indicating the last block where the counter txc_r has been updated, denoted by txb_r, vi) maximum delay guarantee prior to the (on-chain) submission of off-chain transactions to the RSC, denoted by $d_{r,min}$, vii) counter measuring delayed payments, denoted by cdp_r, and viii) original and remaining amount of the penalty fund, denoted by Y_r^o and Y_r, respectively.

2) Relay licensing mechanism: Starting with the *relay licensing epoch* $e-1$ and ending T_R blocks before the new target epoch e, relay licensing is based on a mixed auction and staking process that i) enforces candidate relays to transfer (at least) a minimum penalty fund to the RSC according to the parameters of the requested license and ii) enables RE-CENT nodes to place their stakes as relay witnesses (r-witnesses) and free-of-charge servers in favor of given candidate relay.

TABLE I
EXAMPLE OF TARIFF TABLES (FEES IN RE-CENT COINS - RCS)

	No. of. users	**Fee per user**
Max. no. of Users Tariff Table	1 - 1.000	1.000 RCS
	1.001 - 10.000	1.050 RCS
	10.001-100.000	25 RCS
	100.001-1.000.000	12.5 RCS
	1.000.001-	10 RCS
	No. of. coins	**Fee per coin**
Max. no. of coins Tariff Table	1 - 1.000	0.5 RCS
	1.001 - 10.000	0.2 RCS
	10.001-100.000	0.1 RCS
	100.001-1.000.000	0.01 RCS
	1.000.001-	0.001 RCS
	Tx per block (Tpb)	**Fee per 1 Tpb**
Mean Tx throughput Tariff Table	0 - 0.0001	10.000 RCS
	0.0001 - 0.01	120.000 RCS
	0.01-1	150.000 RCS
	1 - 100	200.000 RCS
	100 -	1.000.000 RCS

Penalty fund calculation and transfer. Each candidate relay uses the tariff tables to calculate the minimum tariff it should pay for the maximum number of relay users it targets to support, the maximum number of user coins it can handle with manageable risk and the number of transactions per block that it will spur into the RE-CENT blockchain. Appropriate selection of these parameters as well as delay and fee values $d_{r,min}$ and o_r, requires deep understanding of how a given relay maps off-chain payments to on-chain transactions. For example, for the given tariff specified in Table I, a candidate relay that requests for a relay license enabling the i) support of up to 10.101 users, ii) management of up to 580.035 coins, and iii) mean throughput of 11.52 transactions per block, should submit to the RSC a minimum penalty fund of 252.525 + 5800.35 + 2.304.000 = 2.562.325,35 RCS.

Auction and staking logic. Candidate relays shall transfer to the RSC the minimum penalty necessary for their license, adding to the transaction payload i) the parameter values defining the license request, ii) the transaction fee $o_r[e]$ and iii) the maximum delay guarantee $d_{r,min}$. To increase the probability of receiving a license, candidate relays may update their offer by assuming higher tariffs. RE-CENT nodes are allowed to time-lock (to the RSC) any portion of their balance in favor of candidate relay(s), acting as *r-witnesses* placing *r-witness stakes*. RE-CENT nodes that to aim to attract more r-witnesses are enabled to provide free-of-charge service offers, by transferring to the RSC the fee of $f_R[e] \cdot X_R[e]$ RE-CENT coins, where $X_R[e]$ is the value of free MBs offered on a per r-witness basis. If the relay receives a license, the RSC shall time-lock penalty funds, r-witness stakes and free-of-charge fees by the end of target epoch. If not, the RSC shall release the respective coins by the end of the licensing epoch.

T_R blocks before the end of the licensing epoch, the RSC will order the list of candidate relays based on the sum of relay penalty fund, r-witness stakes and free-of-charge service fees. License regulation at this step is the key for avoiding poor relay performance due to insufficient transaction capacity of the RE-CENT blockchain. To this end, the RSC will scan the

ordered list from top to bottom and grant license to relays with transaction throughput requests summing up to the estimated Tx throughput $\bar{R}$ (skipping entries violating this criterion). Other criteria can also be used to regulate the number of licenses, e.g. order the list based on the maximum number of users supported per transaction throughput unit. Nonetheless, we consider that the proposed stake-based mechanism will advance system robustness in the long-term.

3) Relay monitoring, reward and penalty mechanisms: RE-CENT nodes may use the relay service by attaching (time-locking) some amount to the RSC in favor of a given (active) relay. The penalty mechanism for free-of-charge RE-CENT servers that do not comply with the promised service offers is assumed similar to the one followed in the VSC contract.

Reward and penalty mechanism for delayed payments. Relays that fully comply with the maximum tolerable delay of attached RE-CENT servers (included in the payload of transactions signed by the RE-CENT client - see Section III.B), shall receive the transaction fees $f_V[e]$ of relevant off-chain payments. However, when relays fail (or refuse) to forward off-chain transactions to the RSC on time, shall trigger a process that will increase the counter cdp_r of delayed payments and will incur an exponentially increasing penalty ${v_R}^{(cdp_r)}$ to both the relay's penalty fund and the user stakes balance (coin burning). Adding to this, the RSC shall directly transfer the respective amount of coins from the balance of the relay user (RE-CENT client) to the RE-CENT server that submits signed transactions proving delayed payments.

If the balance of the relay user (attached to the relay) is not sufficient to support the payment indicated in the delayed payment proofs (i.e. an event indicating that the relay may perform non-legitimate payments), the RSC will ask from the relay to submit proofs of additional funds transferred to the RE-CENT client balance. If the relay fails to do so, its license is revoked and the remaining penalty fund Y_r shall be used to compensate RE-CENT servers submitting proofs of delayed payments. If the residual penalty fund Y_r becomes lower than $x_R Y_r^o$, then the license of relay r is revoked and the RSC will use the all funds, stakes and fees attached to relay r to compensate RE-CENT clients that submit proofs of delayed payments. Provided that Y_r^o is calculated for a given maximum number of attached users, reliable compensation comes down to the careful selection of the x_R and the tariff tables.

Monitoring of the mean transaction throughput. The transaction rate of relays increases when they forward on-chain transactions (calls) to the RSC to update the balance of relay users. The RSC enforces compliance of relays with the throughput agreed in their license, by monitoring its value in short time periods of size k_R blocks (referred to as k-periods). Every time the RSC is triggered by an on-chain transaction by relay r, it increases the counter txc_r by the number of transactions included in the payload of the respective transaction and compares its current value to $k_R x\bar{R}_r$, i.e. the mean number of on-chain transactions allowed within k_R blocks. txp_r is also updated accordingly while txc_r should be assumed zero at the start of a k-period. Note that the RSC will not have to monitor this value of it does not receive calls by the relay (zero throughput). Transactions violating the condition $txc_r \leq_R x\bar{R}_r$ are dropped, triggering delayed payments at the RE-CENT clients. Relays shall avoid such an event to avoid license revoking. Low k_R values enforce strict compliance of relays with their license at the cost of potential underutilization of the transaction capacity of the RE-CENT blockchain (i.e. validators will not increase the block size due to relay throughput control). Large k_R value may result to higher delays in the short-term due to transaction bursts.

Calculation of the mean transaction throughput $\bar{R}$ Parameter $\bar{R}$ defines a stopping criterion during the licensing epoch. It can be calculated based on the transaction throughput agreed in the licenses of the previous epoch and the counters measuring delayed payments. Furthermore, it can be adjusted to (indirectly) increase the transaction capacity for the RE-CENT blockchain. Nonetheless, this approach may trigger unnecessary penalties to honest relays. If relays can be enforced to send honest timestamps to the RSC, the use of relay timestamps can also drive accurate estimation of $\bar{R}$ using simple techniques, e.g. auto regressive moving average.

IV. Numerical Results

In this section, we use analytical expressions derived by [17] to assess the performance gains of the RE-CENT mobile data access model in multi-tier networks co-utilizing multiple spectrum bands. The probabilities of jamming and black hole attacks are set to zero [17]. We consider an eight-tier HWN consisting of two cellular MNOs and three WiFi tiers. The first cellular MNO, tagged as C_1, is composed by tiers 1, 2 and 3 that are considered to include macro, femto and micro base stations (BSs), respectively. Tier-1 macro and tier-2 femto BSs co-utilize band b_1, while tier-3 micro BSs in band b_2. The second cellular MNO, coined as C_2, is composed by tier-4 macro and tier-5 femto BSs that co-utilize band b_3. Tiers 6 and 7 are consider to host Wi-Fi access points (APs) co-utilizing band b_4. Tier-8 Wi-Fi APs operate in band b_5.

We focus on the coverage performance of a typical user, assuming access to tiers 1,2 and 6 through fixed-term contracts. By letting $\mathbb{T}$ denote the set of accessible tiers, the baseline scenario assumes $\mathbb{T} = \{1, 2, 6\}$. The remainder system parameters are set as follows: i) minimum SINR threshold for all tiers is $\gamma = 1.01$ $(0.04dB)$, ii) path loss exponent per band $a_1 = 3.8$, $a_2 = 3.8$, $a_3 = 3.8$, $a_4 = 3.6$ and $a_5 = 3.6$, iii) network intensity per tier $\lambda_1 = 10^{-6}$, $\lambda_2 = 10^{-3.5}$, $\lambda_3 = 10^{-4}$, $\lambda_4 = 10^{-6}$, $\lambda_5 = 10^{-3.5}$, $\lambda_6 = 10^{-2.8}$, $\lambda_7 = 10^{-2.8}$, $\lambda_8 = 10^{-2.8}$, and iv) transmit power per tier $P_1 = 1.5W$, $P_2 = 0.2W$, $P_3 = 0.5W$, $P_4 = 1.5W$, $P_5 = 0.2W$, $P_6 = 0.5W$, $P_7 = 0.5W$, $P_8 = 0.5W$ [17] .

In Fig. 2, we evaluate coverage performance under different combinations of $\mathbb{T}$ that assume access to additional tiers using the RE-CENT access model. As expected, the performance gains (as compared to $\mathbb{T} = \{1, 2, 6)$ vary depending on the number, type and operating band of the additional tiers that can be accessed. Granting additional access to tiers that exclusively utilize a frequency band is preferable ($\mathbb{T} = \{1, 2, 6, 7\}$ vs.

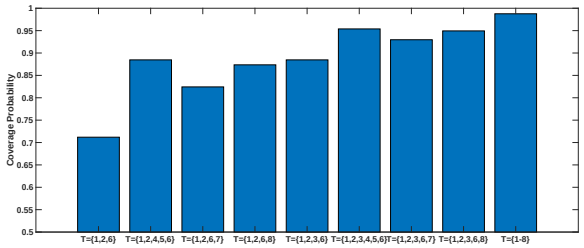

Fig. 2. Coverage probability vs. set of accessible tiers $\mathbb{T}$

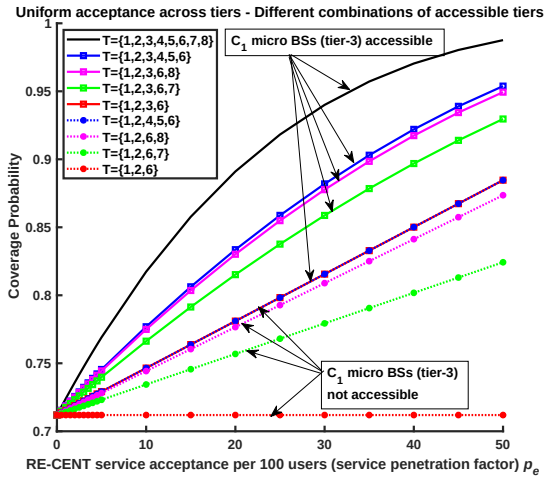

Fig. 3. Coverage probability vs. RE-CENT service penetration

$\mathbb{T} = \{1, 2, 6, 8\}$), while granting additional access to a cellular or a Wi-Fi MNO with exclusive band utilization exhibits similar performance gains for the given model setup parameters ($\mathbb{T} = \{1, 2, 4, 5, 6\}$ vs. $\mathbb{T} = \{1, 2, 6, 8\}$). The type of cellular tiers has little impact on the performance gains attained if no additional access to teirs that operate in a new band is not provided ($\mathbb{T} = \{1, 2, 4, 5, 6\}$ vs. $\mathbb{T} = \{1, 2, 3, 6\}$).

The aforementioned performance trends should be carefully taken into account at both the RE-CENT server, which is required to estimate the promised QoE/QoS level, and the RE-CENT client, which chooses among different RE-CENT servers. Notably, wide acceptance of the RE-CENT paradigm across tiers and end users can maximize network coverage by removing the burden of a priori fixed-term contracts across stakeholders ($\mathbb{T} = \{1, 2, 6\}$ vs. $\mathbb{T} = \{1 - 8\}$).

Fig. 3 plots network coverage for different acceptance levels of the RE-CENT service (measured by the service penetration factor p_e). Results assume uniform acceptance across new tiers: $\lambda_i^* = p_e \cdot \lambda_i$. Even low acceptance of the RE-CENT access model (e.g. an acceptance rate of 30-40% improves network coverage by 10%) can significantly increase the performance experienced by end users ($\mathbb{T} = \{1, 2, 6\}$ vs. $\mathbb{T} = \{1, 2, 6, 7\}$). Besides, rapid increase of performance gains is shown when the user is within coverage of a larger number of RE-CENT enabled tiers, e.g. for $\mathbb{T} = \{1-8\}$, an acceptance rate of 10% proportionally improves network coverage by 10%., highlighting the high potential of RE-CENT access.

V. Conclusions

In this paper we have proposed the RE-CENT mobile data access model enabling dynamic trading of video content on-top of underutilized 5G network assets. Founded on the concepts of DPoS over 5G and off-chain payments using trusted relays, we have overviewed the system architecture and defined the logic of two smart contracts that can implement the RE-CENT access model. Analytical results revealed the high potential of the RE-CENT access model. Future work includes open-source release of the RE-CENT crypto-currency platform as well as meticulous design and implementation of necessary network-level service protocols required to support it.

Acknowledgment

This research has been co-financed by the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call RESEARCH – CREATE – INNOVATE (T1EDK-03524).

References

[1] National Instruments Report on "3GPP Release 15 Overview", 2018.
[2] M. Jiangy, D. Xenakis et al., "Radio Resource Sharing as a Service in 5G: A Software-defined Networking Approach", *Elsevier Comp. Commun.*, vol. 107, no. 15 pp. 13–29, Jul 2017.
[3] S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System", May 2008. [Online]. Available: https://bitcoin.org/bitcoin.pdf
[4] Y. Xiao, N. Zhang, W. Lou, Y. T. Hou, "A Survey of Distributed Consensus Protocols for Blockchain Networks", [Online], https://arxiv.org/abs/1904.04098.
[5] W. Wang et al., "A Survey on Consensus Mechanisms and Mining Strategy Management in Blockchain Networks", *IEEE Access*, vol. 7, pp. 22328-22370, 2019.
[6] I. Barinov et al., "POSDAO: Proof of stake decentralized autonomous organization", *Technical report*, SSRN, Sept 2019. [Online] http://dx.doi.org/10.2139/ssrn.3368483.
[7] Parity Documentation - Authority Round. http://wiki.parity.io/Aura.html.
[8] Ericsson, "Ericsson Mobility Report", *Technical Report*, June 2019.
[9] D. Rusako et al., "First Internet Marketplace power by P2P VPN Network on Blockchain", *White paper*, May 2019. [Online]. https://www.cryptoground.com/privatix-white-paper
[10] D. Tsolkas et al., "A survey on parametric QoE estimation for popular services," *Elsevier J. of Net. and Comp. App.*, vol. 77, pp. 1-17, 2017.
[11] Z. Xiong et al., "Cloud/Fog Computing Resource Management and Pricing for Blockchain Networks", *IEEE Internet of Things*, vol. 6, no. 3, pp. 4585-4600, June 2019.
[12] D. Xenakis et al., "Mobility Management for Femtocells in LTE-Advanced: Key Aspects and Survey of Handover Decision Algorithms", *IEEE Surv. and Tut.*, vol.16, no.1, pp.64-91, Q1 2014.
[13] E. Liotou et al., "QoE-SDN APP: A rate-guided QoE-aware SDN-APP for HTTP adaptive video streaming", *IEEE J on Sel. Areas in Commun.*, vol. 36, no. 3, pp. 598-615, March 2018.
[14] G. Wood, "Ethereum: a secure decentralised generalised transaction ledger", Byzantium Version, Aug. 2019. [Online] https://ethereum.github.io/yellowpaper/paper.pdf
[15] J. Poon, T. Dryja, "The Bitcoin Lightning Network: Scalable Off-Chain Instant Payments", *White Paper*, 2016. [Online] https://lightning.network/lightning-network-paper.pdf
[16] A.E. Gencer et al., "Decentralization in Bitcoin and Ethereum Networks", *Int. Conf. on Fin. Crypt. and Data Sec.*, pp 439-457, Dec 2018.
[17] A. Tsiota, D. Xenakis et al., "On Jamming and Black Hole Attacks in Heterogeneous Wireless Networks", *IEEE Trans. on Veh. Tec.*, to appear.