

Proof of Networking: Can Blockchains Boost the Next Generation of Distributed Networks?

Lorenzo Ghio, Leonardo Maccari, Renato Lo Cigno
Dept. of Information Engineering and Computer Science, University of Trento, Italy
{name.surname}@unitn.it

Abstract—The recent explosion of interest in blockchains led to a plethora of proposals for their application, including attempts to decentralize some centralized network functions. At the same time, real “distributed wireless networks” are emerging. Community networks, for instance, are large mesh networks made of hundreds of nodes built by communities primarily to solve digital divide, and they are thriving. The challenges these networks face are not only technological: they deal with creating incentives to participate, with the business model they may adopt, and with their internal governance. Very few models have been proposed to apply blockchains to bottom-up distributed networks: we instead expose how they can solve many problems which so far hindered the diffusion of such networks. Maybe we can push this further: a network is, in essence, a system in which all nodes find a rough consensus on the best paths to connect a node with another. Can we use this consensus method to run a distributed ledger and a cryptocurrency within the network itself, rather than simply applying to networks the effects of a blockchain defined in a separate system? This paper introduces this concept, named “Proof of Networking”, and discusses its potential avails.

I. INTRODUCTION

Wireless Mesh Networks (WMNs) have been one of the main staples of both networking and distributed systems research, with visionary works imagining that people would use ad-hoc networks to communicate with their peers, and cables would disappear replaced by WMNs. But actually, in our daily life, we communicate with our peers using the cellular infrastructure and the vast majority of households are connected with cables to a global wired infrastructure. In this paper we argue that the main reasons for this situation are not technological but somehow *environmental*, and we describe some of the obstacles that hindered their diffusion asking: Can blockchains remove these obstacles?

To address this question, we analyze the typical use cases for WMNs, we review the technology underpinning blockchains and a few noteworthy applications in the networking domain, next we mix the two concepts. We argue the advantages of porting blockchains into distributed networks (WMN being the most prominent technology), and we finally introduce the concept of “Proof of Networking” as a full, enabling merge for distributed, on-demand networking.

This research is supported by the European Commission, H2020-ICT-2015 Programme, Grant Number 688768 ‘netCommons’ (Network Infrastructure as Commons).

II. AD-HOC AND MESH NETWORKS

Ad-hoc and mesh networks are both characterized by multi-hop communications among devices composing a network that is unplanned, dynamic and self-healing. They mainly differ in scale and use.

A. Ad Hoc Networks

An ad-hoc network is in general imagined as a strictly local network made of a few (tens) portable devices, with some degree of mobility.

Android and iOS devices alone are more than 3 billions. Almost each of these is equipped with a modern Wi-Fi chip, and thus would virtually enable the creation of ad-hoc networks. However, this does not happen for a simple reason: The Operating System (OS) does not allow it.

It’s a deliberate choice to disable this feature: the kernels at the core of Android and iOS support it. Ad-hoc networks may be technically feasible, but the OSs prevent us from making an attempt at them, looking for a profitable use case. Emblematic is the “bug 82”, opened in the Android tracker by a user asking for ad-hoc mode support in 2008: Its current status is “Won’t Fix”. Anyone can speculate on the reason, but the ground truth is that there is no technical reason to avert it, thus it is a business (or policy) choice to restrain ad-hoc networking, that we can imagine can/will be removed if there is a valid case to do so.

B. Mesh Networks

A mesh network is instead imagined as a mainly static network made of wireless nodes covering areas that range from a house to a whole city. Mesh networks had better luck in the last decades, they are used in both industrial and military applications, and lately also to extend indoor WiFi coverage. The most notable application of mesh networks are in all likelihood Community Networks (CNs). A CN is a (wireless) mesh created by a community of people, primarily to solve a condition of digital divide. The concept is not new, but their growth in the past few years has been remarkable [1]. CNs drew the attention of several research disciplines because they are bottom-up socio-technical experiments of networking that can alleviate the digital divide [2], yet their growth and existence is always in jeopardy, just like ad-hoc networks.

Regardless of ad-hoc or CN flavor, WMNs have a meaning if they serve a human community. A human community, below a certain size, is small enough to self-organize without

any formal agreement. Beyond that size (around tens of participants [3]) informal organization does not work anymore, and also collusion may arise [4]. A thorough study in the more structured CNs [5] outlined negative behaviours that emerge regularly, we specifically consider two of them: the “dumping” and the “club of techies” problems.

The first concerns, in a voluntary system, the struggle to nurture participation of people to maintain the system alive, which requires acknowledgement of the system’s value. But the value of a network is hard to perceive: users perceive value in the applications they use, but consider the network as a commodity, something they can give for granted. Consequently, many people participating to CNs stop managing their node right after they turn it on, and let it degrade, as long as this does not directly influence their own experience.

The second pattern describes the case in which a small number of tech-savvy people start building the network as a voluntary effort to solve their own problems. The network grows and the core group of maintainers can not cope with the effort needed to manage it anymore. Since the “club of techies” was never interested in delegating responsibilities to other (possibly non-techie) people, the network collapses.

These phenomena are well known, but not solved, in social economics, and the study of Common Pool Resources (CPRs) gained the Nobel prize to Elinor Ostrom. A CPRs is a shared good governed by a local community with an internal and mostly horizontal governance. In situations in which both the market and state intervention fail to efficiently manage a certain resource, CPRs have been shown to be effective. In the context of distributed networked systems, it is clear that the network itself is the CPR and it is necessary to find a technically feasible and economically viable way to maintain it so that applications, where a market can be established, can flourish. It must be noted, however, that most successful WMNs operate in situations of market failure. They exist thanks to voluntary work of people which reduces costs of ownership of the whole infrastructure. Although not mandatory, a governance scheme for CNs must blend both the economic and social incentives which drive such networks.

The question we pose at this point is: Can blockchains be of use to solve these problems? Before giving an answer, let us recall briefly what a blockchain is and how it can be used in a network.

III. BACKGROUND ON BLOCKCHAINS

Albeit known before, blockchains became famous after Nakamoto¹ based the first successful crypto-currency system on a blockchain technology, the Bitcoin [6], where the main problem is to ensure validity of transactions (TXs) in the absence of a central authority, preventing the “double spending” of virtual currency. This is not trivial because Bitcoin nodes use a P2P network to publish their TXs and, due to propagation delays, validators may receive two distinct TXs that spend the

¹Satoshi Nakamoto is the pseudonym used by the author, or authors, of the white paper that unveiled the Bitcoin system in 2008.

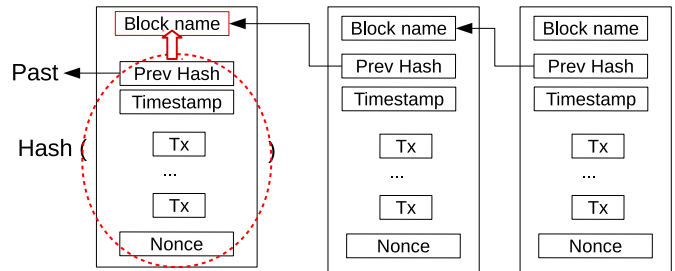


Figure 1: Blocks are linked via hash-pointer to previous block

same money in different order. They therefore need to find consensus on which came before, and is valid, and which after, and is not. Nakamoto proposed to group TXs in a chain of timestamped blocks (blockchain), thus implementing a TXs ordering system together with a consensus mechanism. Each block contains a set of TXs, a time stamp, the hash fingerprint of the previous valid block (i.e. its “name”), and a nonce, as shown in Fig. 1; A block is valid if its content, hashed with a double SHA256, produces a fingerprint that has a pre-defined number of leading zeros. The node that generates a block is also allowed to add a TX to itself of a certain amount of Bitcoin that were “mined” in this process, as a reward for its work. A blockchain is immutable meaning that, once a block is added, it can not be removed or corrupted. In fact, every modification of a past block would invalidate its hash name, which in turn invalidates all the subsequent blocks. Finding a valid nonce is computationally intensive, and this is why this system is called Proof of Work (PoW). All miners compete to find the next block, which leads to the well known energy consumption problem of Bitcoin.

Basically a PoW-based blockchain is a data structure that implements a distributed and tamper-proof Shared Ledger (SL) in a *trustless* network. PoW is not the only way to reach consensus on new chain’s blocks, more energy efficient methods exist, especially if the network is not fully trustless. A review of blockchain technology is out of the scope of this paper, and several tutorials are available [7].

A. Blockchains in Networks

A multitude of approaches exist to take advantage of blockchains in networked applications, but only a few of them push blockchains down to the network level. Among them we mention a technique to jointly address contractual and routing issues typical of interdomain routing [8] and a mechanism to speed up the synchronization of consumers’ status-updates in a Named Data Network [9]. Sharing of blacklists between ASs for security applications has been proposed [10], and applications in the IoT domain [7] as well. Several projects try to decentralize some of the still centralized network functions, such as DNS². Finally, two relevant projects for our context are Althea and AMMBR³, which try to implement a sustainable distributed wireless network, both relying on blockchains.

²<https://emergoin.com>

³<http://ammbr.com>, <http://altheamesh.com>

The common ground of both projects is the creation of a distributed marketplace to incentive participation and competition among those who want to act as service providers. To this end, they prompt microtransactions in cryptocurrency either to finance local services or to pay for Internet access.

Althea incentives peering agreements between nodes. Node A, with Internet access, will broadcast in beacons the expected quality using the ETX metric and its price per kilobyte. Node B, that needs Internet access, will do a peering agreement with A using a micro-transaction before actually establishing a working link at the IP layer. Node B may re-sell access to a third node C that does not have direct visibility of node A, and so on.

Althea nodes do not mine blocks, they rather use an external blockchain (the Ethereum blockchain) with Micropayment Channels⁴. In short neighbouring nodes, before creating a link, pre-charge some credit with an empty transaction on the blockchain. This means that they agree on the Ethereum price to forward given amounts of bytes, then start performing frequent local TXs. The local link-balance, frequently updated off-chain, is only infrequently synced with the blockchain. Despite the efficient payment system, the proposed metering mechanism to verify the service-level and, consequently, authorize or deny payments, is to be further studied and verified.

The AMMBR blockchain instead is a dedicated immutable ledger to record pricing, metering, billing, payment, reconciliation, reporting and auditing. In practice, the blockchain is used for peering but also to support the presence of services inside the network, which can be acquired via the blockchain, and can use the blockchain for various tasks, for instance, identity management. AMMBR uses a proprietary chip to replace PoW with the *proof of Elapsed time* (PoET). The PoET workflow implements a fair and random leader-election algorithm but seems to be depending on a centralized server that plays a key coordination role. While Althea already published a detailed white paper and some open source code, the details of AMMBR are still not public. However, AMMBR representatives committed to publish open source code and design, so we will have to wait for their first specifications.

IV. PROOF OF NETWORKING

AMMBR and Althea introduce blockchains and transactions to foster an in-network marketplace. This can make local networks more appealing, as long as they enable Internet access and local services as we discuss in Section IV-A, but the blockchain and the network remain separate entities.

A. Enabling a Marketplace

Imagine a group of people with their mobile devices; a bus carrying the group to fix ideas. Imagine that some of them do have Internet access, and that an application activates the WiFi interface and announces in beacons the offer of Internet access at a certain price. Imagine now that this can happen, as never

⁴A payment channel is a trusted method for two parties to exchange payments by signing transactions that alter the balance of an escrow account held by a bank or blockchain

before, in a multi-hop ad-hoc network in which each node pays for the access and may even resell it to some other neighbor. Such a system requires fast micro transactions without fees to take place between people that do not trust each other, and a cryptocurrency fit for the purpose. We have created a blockchain enabled marketplace on top of an ad-hoc network, and now we only need to exert our fantasy.

We can extend the scarcity problem, where most devices do not have Internet access, to a quality problem: most devices have Internet access but they have different quality (they use different operators or different generations of the wireless technology). Mobile devices could seamlessly choose which is the option that better suits their current need and perform a micro transaction to enable it. Similarly, devices could offer local services such as media sharing, proxy access to external services, on-line gaming and caches of software updates, and use transactions to enable them.

These visions are not new, but never materialized because they never had a working business model behind them that could make OS vendors change their mind and enable ad-hoc networks on their platform.

B. Embedding the Blockchain

The blockchain, in the model described so far, is just an enabler of transactions, and it can reside outside the network itself. In other terms, it uses the network for communications but the transaction recording and the “proof” for them are independent from the network. The network only partly benefits from the presence of a blockchain, it mostly benefits from the presence of a currency and a way to inject it into the network. Can we instead embed the blockchain into the network?

Let us recall that, in practice, a blockchain is a method to obtain consensus leveraging some sort of “proof”: of work, of time elapsed, or any other proposed one. And what is an IP network? It is a system in which nodes need to reach a consensus on the way to go from node A to node B, and the proof is the delivery of packets: We have obtained a Proof of Networking (PoN). If we take for instance a link-state routing protocol, it distributes information so that all nodes share the same view of the network graph. If this does not happen the network simply does not work, e.g., routing loops are created.

Consider a mesh network, in which node A uses a link-state protocol with embedded cryptographic signatures, on the model of Secure-OLSR [11]. Node A performs link-sensing with its neighbors and periodically floods signed TC messages (containing its active links) to the whole network. At network convergence every node should have enough information to know the whole graph topology, with the information on each link being cross signed by both endpoints.

If this information is periodically “frozen” and agreed upon, it provides the “proof of networking” needed to quantify the value of the network and the contribution of individual nodes. The required frequency depends on the scenario: once per day may be enough for a stable community network, while small ad-hoc networks built on the fly may require a much smaller interval. The proof can be built selecting one node, for instance

the one with highest centrality (several centrality metrics exist to rank nodes in a graph) that broadcasts a message containing the topology made of all the cross-signed links. The message is flooded to the network, travelling only on links that are not only existing but also present in the frozen topology. A block can be generated including this proof together with additional transactions. Some of them could be used to generate new currency based on the value of the network (like in the Bitcoin scheme), while other ones could settle the services generated by nodes in the last interval to quantify, e.g., the traffic routed by each node or other services provided to participants.

This approach embeds the blockchain and the token-currency into the network, with direct benefits leveraging on the “network effect” and all the advantages induced by an integrated ecosystem. Let’s analyze some possible features of a tokenized system:

- a) The amount of generated currency can algorithmically depend on the network evolution (as in the Bitcoin, where it depends on the total computing power), which fosters the network growth;
- b) Rewards can be delivered to all nodes proportionally to their importance in the network topology, and on the traffic they carry. This way nodes are incentivized to be central in the topology, thus routing a high quantity of traffic, and not remain leaf “free rider” nodes;
- c) Rewards should acknowledge collective behaviours, e.g. introducing a dependency between the number of tokens generated and the number of nodes added in past intervals;
- d) The network graph can be enriched with annotations that include other network parameters (forwarded traffic, uptime, list of supported services, ...) to represent a composite metric of contribution to the network value in multiple dimensions.

This model can help solving the governance problems we mentioned in Section II-B. The value of the network infrastructure is finally quantified: It resides in the importance of nodes and in their contributions. If one does not put enough effort in the maintenance of his node, the whole network will lose some value and he will be tangibly affected. Conversely, maintaining a central node efficient rewards the owner with currency to be spent for access to Internet or other services, but requires effort to upgrade the node and make it work properly. The “club of techies” that bootstrap the network will have an incentive in involving more maintainers, and the network growth will be shared and fairly distributed.

Currency does not necessarily mean real money. Virtual tokens, or any other currency used by a local community, can be used to quantify the work needed to maintain a node, and the value of services received. This would preserve also the voluntary-based approach of many CNs.

At a larger scale things can get even more interesting. Different networks will use different currencies, requiring gateways as brokers to exchange different currencies, this to perform automatic peering agreements between different communities that can not physically merge in just one network. This would let mesh networks scale up to their physical limit connecting them with different technologies. While this happens normally

with ISPs on a commercial ground, a blockchain/cryptocurrency approach would simplify the acknowledgement of internal value in different systems, thus incentive creation of links and connections between heterogeneous networks.

V. TECHNICAL CHALLENGES & THE WORK AHEAD

Creating such a “proof of networking” is technically challenging, we can foresee at least three themes that must be carefully analysed.

The first is the specification of the PoN, of which we gave only a sketch. In fact, a network needs an “approximate consensus” to work, meaning that when the network changes, e.g. a link is removed, there is a transitory phase in which nodes may disagree on some pieces of information. While we accept potential service degradation before convergence of a routing protocol, degrading the security level of a system, even temporarily, could be catastrophic. To counter this, there exists a large body of research on mesh/ad hoc network security that can be exploited to meet this goal.

A second theme is the size of the blockchain and of the amount of information to be flooded in the network. This is not a new problem, as the issues with scalability and throughput in blockchains are known ones. While improvements are emerging [12], their applicability to wireless routers needs to be tested.

A third theme is the integration with routing. How do we ensure loop-free and stable routing when the routing metric does not only take into account the link quality, but also embeds the price of the transit for that link? And how can a node be sure that the service-level agreement that it negotiates with a peer is effectively enforced? Here we enter into an innovative research area that rises very stimulating challenges in networking, optimization and economics.

REFERENCES

- [1] L. Maccari and R. Lo Cigno, “A week in the life of three large Wireless Community Networks,” *Ad Hoc Networks*, vol. 24, Part B, Jan. 2015.
- [2] R. Lo Cigno and L. Maccari, “Urban Wireless Community Networks: Challenges and Solutions for Smart City Communications,” in *ACM Int. Workshop on Wireless and Mobile Technologies for Smart Cities (WiMobCity '14)*, part of *MobiHoc 2014*, Aug. 2014, pp. 49–54.
- [3] L. Maccari, “On the Technical and Social Structure of Community Networks,” in *The IFIP Internet of People Workshop (IoP)*, July 2016.
- [4] G. Ciccarelli and R. Lo Cigno, “Collusion in Peer-to-Peer Systems,” *Computer Networks*, vol. 55, no. 15, pp. 3517–3532, Oct. 2011.
- [5] S. Crabu, L. Navarro, M. Dulong de Rosnay, D. Franquesa, and F. Tréguer, “Report on the Governance Instruments and their Application to CNs (v1), netCommons Deliverable D1.3,” June 2017.
- [6] S. Nakamoto, “Bitcoin: A peer-to-peer electronic cash system,” 2008.
- [7] K. Christidis and M. Devetsikiotis, “Blockchains and Smart Contracts for the Internet of Things,” *IEEE Access*, vol. 4, May 2016.
- [8] I. Castro, A. Panda, B. Raghavan, S. Shenker, and S. Gorinsky, “Route Bazaar: Automatic Interdomain Contract Negotiation,” in *15th Workshop on Hot Topics in Operating Systems*, May 2015.
- [9] T. Jin, X. Zhang, Y. Liu, and K. Lei, “BlockNDN: A bitcoin blockchain decentralized system over named data networking,” in *2017 International Conference on Ubiquitous and Future Networks (ICUFN)*, Jul. 2017.
- [10] B. B. Rodrigues, T. Bocek, and B. Stiller, “Multi-domain DDoS Mitigation Based on Blockchains,” in *IFIP International Conference on Autonomous Infrastructure, Management and Security*, Jun. 2017.
- [11] H. Fan, H. Liang, and F. Cai, “Secure OLSR,” in *19th Int. Conf. on Advanced Information Networking and Applications*, March 2005.
- [12] I. Eyal, A. E. Gencer, E. G. Sirer, and R. Van Renesse, “Bitcoin-ng: A scalable blockchain protocol,” in *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation*, 2016.