

Mining competition in a multi-cryptocurrency ecosystem at the network edge: a congestion game approach

Eitan Altman
INRIA and Bell Labs Joint Lab
France

Mandar Datar
INRIA and Bell Labs Joint Lab
France

Alexandre Reiffers
Indian Institute of Science
Bangalore, India

Swapnil Dhamal
INRIA
France

Daniel S. Menasché
UFRJ
Rio de Janeiro, Brazil

Corinne Touati
INRIA
France

ABSTRACT

We model the competition over several blockchains characterizing multiple cryptocurrencies as a non-cooperative game. Then, we specialize our results to two instances of the general game, showing properties of the Nash equilibrium. In particular, leveraging results about congestion games, we establish the existence of pure Nash equilibria and provide efficient algorithms for finding such equilibria.

1. INTRODUCTION

The blockchain is a distributed synchronized secure database containing validated blocks of transactions. A block is validated by special nodes called miners and the validation of each new block is done via the solution of a computationally difficult problem, which is called the proof-of-work puzzle. The miners compete against each other and the first to solve the problem announces it, the block is then verified by the majority of miners in this network, trying to reach consensus. After the propagated block reaches the consensus, it is successfully added to the distributed database. The miner who found the solution receives a reward either in the form of cryptocurrencies or in the form of a transaction reward.

Because of the huge energy requirement necessary to be the first to solve the puzzle, blockchain mining is typically executed in specialized hardware. In [20] an Edge computing Service Provider (ESP) is introduced to support proof-of-work puzzle offloading by using its edge computing nodes. In [19] a game is formulated between the miners in the presence of a single ESP and then a Stackelberg game is used to compute the pricing that maximizes the revenue of the ESP.

Our work addresses the following two questions:

- 1) given a single blockchain, how should rational users contribute to the mining process, possibly counting on third-party ESPs or mining pools to offload infrastructure costs?
- 2) given multiple blockchains, e.g., in a multi-cryptocurrency ecosystem, how should rational users distribute their monetary and/or computational budget towards mining?

In this note, we focus on the competition between miners while addressing the two questions above. We model the competition over several ESPs and over several blockchains characterizing multiple cryptocurrencies as a non-cooperative game. Then, we specialize our results to two instances of the general game, showing properties of the Nash equilibrium.

In the first game, there is a single blockchain (e.g., cryptocurrency) and any of the M ESPs (or mining pools) can be used by the miners to solve the puzzle. In the second game, we consider K opportunities, each of which corresponding to another blockchain. At each time slot of duration T (which corresponds to a new puzzle to be solved) each of the miners decides which of K puzzles to solve. We formulate both games and establish conditions for the existence of a pure Nash equilibrium for the association problem between miners and ESPs, providing an efficient algorithm for solving it. We summarize our contributions as follows:

Congestion game for mining competition: we model the competition among users searching for a solution to the mining puzzle as a game (Section 2). In essence, as the number of users willing to mine increases, the chances that a given user is the first to succeed in solving the mining puzzle and wins a reward decreases (i.e., the system becomes *congested*). In particular, we assume that users can count on third-parties to offload infrastructure costs, and can mine multiple cryptocurrencies. Under the assumption that such third-parties are roughly indistinguishable, we further show that when there is one single cryptocurrency of interest the *congestion game* admits a simple equilibrium accounting for users that must decide whether to mine or otherwise not join the system (Section 3).

Analysis of multi-cryptocurrency ecosystem: we analyze the congestion game involving multiple cryptocurrencies. In that case, miners compete against those that decide to mine the same cryptocurrency (Section 4). We show that the proposed game admits a potential, and discuss a number of extensions, such as accounting for dynamic puzzle complexity and mining pools (Sections 5 and 6).

2. BLOCKCHAIN COMPETITION GAME

We consider a population of M ESPs and a set of K blockchains. We assume that in ESP m , the amount of service $R_{k,m,i}$ requested by miner i to solve puzzle k is exponentially distributed with expectation $1/\mu_{k,m}$. $R_{k,m,i}$ are independent RVs. Thus, if there are $\ell_{k,m}$ miners associated to ESP m mining currency k , the time it takes for the fastest of them to solve the puzzle corresponding to currency k is exponentially distributed with expectation $1/(\sum_m \mu_{k,m} \ell_{k,m})$.

Miners, mining servers and puzzles. We denote by $\mathcal{N} := \{1, 2, \dots, N\}$ the set of miners. There is a finite population of miners, and if a miner changes his strategy this will cause a change in the utilities of other miners. Let

$\mathcal{K} := \{1, 2, \dots, K\}$ be the set of puzzles, each of which associated with a different cryptocurrency that the miners are trying to solve. We assume that each cryptocurrency corresponds to exactly one puzzle. Let $\mathcal{M} := \{1, 2, \dots, M\}$ denote the ESPs, also referred to as mining servers, that miners can rely on. Notation is summarized in Table 1.

Strategies. Set $\mathcal{S}_i \subset \mathcal{K} \times \mathcal{M}$ denotes the set of ordered pairs (puzzle, ESP), corresponding to ESPs that miner i can rely on to solve puzzles of a given type. The set \mathcal{S}_i can differ across miners due to political or economic restrictions. For instance, certain countries do not allow investment in certain cryptocurrencies. Alternatively, the set of available ISPs for two different miners may not be the same. A strategy for miner i is denoted by $s_i \in \mathcal{S}_i$, corresponding to the puzzle (cryptocurrency) which the miner intends to solve using a given infrastructure. A strategy vector $s := (s_i)_{i \in \mathcal{N}}$ produces a load vector $\ell := (\ell_{k,m})_{k,m}$.

Rewards, costs and utilities. Let η_k be the load of miners across all ESPs towards cryptocurrency k . Then,

$$\eta_k := \sum_{m' \in \mathcal{M}} \ell_{k,m'} \mu_{k,m'}. \quad (1)$$

Recall that for a given load vector ℓ , the time to solve the puzzle of the k^{th} cryptocurrency is exponentially distributed with expectation $1/\eta_k$. Let q_k be the probability that puzzle k is solved by time T ,

$$q_k = 1 - \exp(-T\eta_k). \quad (2)$$

Let $\tilde{p}_{k,m}$ denote the probability that a miner using ESP m is the first to solve puzzle k . Then, $\tilde{p}_{k,m} = q_k \ell_{k,m} \mu_{k,m} / \eta_k$. Throughout this paper, $0/0 = 0$. In the expression of $\tilde{p}_{k,m}$, for instance, if $\eta_k = 0$ and $\ell_{k,m} = 0$, then $\tilde{p}_{k,m} = 0/0 = 0$.

The probability that a given miner using ESP m is the first to solve puzzle k is

$$p_{k,m} = 1_{\ell_{k,m} > 0} q_k \mu_{k,m} / \eta_k, \quad (3)$$

where 1_c equals 1 if condition c holds and 0 otherwise.

Let $U_{k,m}(\ell)$ denote the utility to a miner who tries to find the solution of the current puzzle associated to cryptocurrency k , using ESP m . The utility is the weighted sum between the probability to solve first the puzzle and the cost associated with it. We denote by $\gamma_{k,m}$ the cost of mining blockchain k at ESP m . Thus,

$$U_{k,m}(\ell) = \begin{cases} p_{k,m} - \gamma_{k,m} & \text{if } p_{k,m} > \gamma_{k,m}, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

User i utility is $\tilde{U}_i(s_i, s_{-i}) = \sum_{(k,m) \in \mathcal{S}_i} 1_{s_i=(k,m)} U_{k,m}(\ell)$, where $s_{-i} := (s_1, s_2, \dots, s_{i-1}, s_{i+1}, \dots, s_N)$ is the vector of strategies of all miners except miner i . Given the ingredients above, the blockchain competition game is characterized by $C = \langle \mathcal{N}, \mathcal{K} \times \mathcal{M}, (\mathcal{S}_i)_{i \in \mathcal{N}}, (U_{k,m})_{(k,m) \in \mathcal{K} \times \mathcal{M}} \rangle$. In Sections 3 and 4 we analyze two special instances of this game.

2.1 Congestion games and potentials

Next, we briefly introduce some basic background on congestion games, crowding games and potentials. Such background is instrumental in the analysis of the blockchain competition game that follows.

Congestion games [12] are equivalent to routing over an arbitrary graph, when all routed objects have the same size, and are non splittable. The cost of using an edge is the same for all players. Crowding games [10] are congestion games

Table 1: Table of notation

variable	description
$\ell_{k,m}$	number of users mining blockchain k at ESP m
$p_{k,m}(\ell)$	probability that user is first to mine a block
$U_{k,m}(\ell)$	utility of user mining blockchain k at ESP m
$\gamma_{k,m}$	mining cost associated to blockchain k at ESP m

with more restricted topology (parallel links) but more general costs (user dependent).

In our setup, the routed object is the mining power. In the simplest setup, wherein we have one single cryptocurrency, the vertices of the graph are users and ESPs (with a virtual ESP corresponding to the option of not mining). The cost incurred by a user who decides to mine through a given ESP is the cost of an edge between the user and the ESP.

A congestion game without player specific payoff functions is guaranteed to admit a standard potential and a pure equilibrium [12]. A game that does not admit a standard potential may still admit an ordinal potential. A game with an ordinal potential can have any finite subset of actions available to a player, still admitting a pure equilibrium.

Milchtaich [10] proves the existence of a pure Nash equilibrium given user dependent costs in crowding games. In this paper, we are interested in user dependent strategy sets. Nonetheless, one can show an equivalence between user dependent costs and user dependent strategy sets, and henceforth we use interchangeably the two notions.

3. ESP ASSOCIATION GAME

In this section, we introduce the ESP association game and analyze some properties of its equilibria.

Next, we consider the special case where we have only one cryptocurrency, which we denote by \star . Given vector ℓ , where $\ell_{\star,m}$ denotes the number of miners associated to ESP m , the time $T^{(1)}$ when the fastest solves the puzzle is exponentially distributed with expectation of $E(T^{(1)}) = 1/\eta_\star$, where $\eta_\star := \sum_{m' \in \mathcal{M}} \ell_{\star,m'} \mu_{\star,m'}$. Note that if $\ell_{\star,m'} = 0$ for all m' then $\eta_\star = 0$ and $E(T^{(1)}) = \infty$. In Section 5 we discuss heuristic strategies adopted by existing cryptocurrencies to prevent the time to solve a puzzle to grow unboundedly.

The probability that a given miner associated with ESP m wins is given by (3), $p_{\star,m} = 1_{\ell_{\star,m} > 0} q_\star \mu_{\star,m} / \eta_\star$ where q_\star is given by (2).

In what follows, we analyze the ESP game proposed above. We assume that a miner has always the option of not associating to any ESP and in that case its utility is zero.

THEOREM 1 (EXISTENCE). *A pure Nash equilibrium exists. We denote it by $s^* := (s_i^*)_{i \in \mathcal{N}}$.*

PROOF. First, consider the case where for all i and j , $\mathcal{S}_i = \mathcal{S}_j$ and s_i does not depend on i . Then, the ESP association game is a congestion game, in the sense of [14], and the theorem follows. Otherwise, the game is a crowding game, in the sense of [10]. As shown in [10], as the game has player-specific payoffs it may not admit a standard potential [12], but it still admits pure Nash equilibria. \square

Let ℓ_\star be the number of miners that decide to associate to an ESP,

$$\ell_\star = \sum_{m=1}^M \sum_{i \in \mathcal{N}} 1_{s_i^*=(\star,m)}. \quad (5)$$

Then, $N - \ell_*$ is the number of users that decide not to mine.

When all $\mu_{*,m}$ are equal we denote them by μ_* . Then, equation (3) reduces to

$$p_*(\ell_*) = 1_{\ell_* > 0} (1 - \exp(-T\mu_*\ell_*)) / \ell_* \quad (6)$$

where p_* is the probability that a user that decides to connect to an ESP is the first to solve the puzzle. Let $\gamma_{*,m}$ be a fixed cost for associating with ESP m .

In the following theorem, we assume that $\mu_{*,m}$ and $\gamma_{*,m}$ are the same for all m (therefore $\mu_{*,m} = \mu_*$ and $\gamma_{*,m} = \gamma_*$ for all m). Let the utility for a miner associating to ESP m be given by (4).

THEOREM 2 (NO PLAYER-SPECIFIC STRATEGIES). *If for all i and j , $\mathcal{S}_i = \mathcal{S}_j$ and s_i does not depend on i , the Nash equilibrium is given by the solution of the following optimization problem,*

$$\operatorname{argmin}_s \Phi(s) := \sum_{l=1}^{\ell_*} p_*(l) - \gamma_* \quad (7)$$

$$\text{subject to:} \quad \ell_* \leq N, \quad \ell_* \geq 0. \quad (8)$$

Equation (11) is the game potential function. The optimization problem (11)-(8) is equivalent to a bin-packing problem with concave costs.

PROOF. This is a congestion game in the sense of Rosenthal [14] and therefore has a potential. Indeed, in this game each player can decide to associate or not with an ESP. Thus all associations to the M ESPs can be aggregated to a single route that represents the choice of mining and the option of not associating represents the second route. \square

THEOREM 3 (PLAYER-SPECIFIC STRATEGIES). *If s_i depends on the identity of user i , the game may not admit a standard potential, but still admits pure Nash equilibria.*

PROOF. The game is a crowding game, and the result follows from [10, 11]. \square

4. BLOCKCHAIN ASSOCIATION GAME

In this section, we introduce the multiple cryptocurrencies game and derive structural properties of the associated set of equilibria.

Here we assume that there are K cryptocurrencies. To simplify presentation, we consider a single ESP, and drop subscript m from all variables.

For a given load vector ℓ , the time it takes till the fastest puzzle to be solved is exponentially distributed with expectation $1/(\mu_k \ell_k)$. Thus, the probability that a miner is the first to solve the puzzle is

$$p_k(\ell_k) = (1 - \exp(-T\mu_k \ell_k)) / \ell_k \quad (9)$$

Note that $p_k = 0$ if $\ell_k = 0$ (recall that we assume $0/0 = 0$ throughout this paper). The utility of a tagged miner to mine a cryptocurrency k when there are ℓ_k miners associated with the same cryptocurrency (including the tagged miner) is given by (4), where

$$U_k(\ell_k) = p_k - \gamma_k, \text{ if } p_k - \gamma_k \geq 0. \quad (10)$$

We add to it the constraint that a miner does not participate in solving the puzzle if its utility is negative. In that case the equilibrium is characterized by the condition

$\sum_k \ell_k^* \leq N$, $\ell_k^* \geq 0$, for $k = 1, \dots, K$. This game is referred to as an *elastic game*. If the equilibrium vector ℓ^* saturates the constraint ($\sum_k \ell_k^* = N$, $\ell_k^* \geq 0$, $k = 1, \dots, K$) then for each k for which $\ell_k^* > 0$, and each k' , $U_k(\ell_k^* - 1) \geq U_{k'}(\ell_{k'}^* + 1)$.

Similar theorems as those presented in the previous section establishing the existence of pure Nash equilibria and characterizing the equilibria still hold under the blockchain association game. The statements of the theorems and the proofs are similar to those in the previous section, and are omitted for conciseness. Recall that in Theorem 2 for all i and j , $\mathcal{S}_i = \mathcal{S}_j$ and s_i does not depend on i . Then, in this case the number of miners associated to each cryptocurrency $\ell_k^* = \sum_{i \in \mathcal{N}} 1_{s_i^* = k}$ is now the solution of the following optimization problem,

$$\operatorname{argmin}_s \Phi(s) := \sum_{k \in \mathcal{K}} \sum_{l=1}^{\ell_k} p_k(l) - \gamma_k \quad (11)$$

$$\text{subject to:} \quad \sum_{k \in \mathcal{K}} \ell_k \leq N, \quad \ell_k \geq 0. \quad (12)$$

Theorem 2 holds replacing (7)-(8) by the equations above.

5. DISCUSSION

Positive and negative externalities. In the models proposed in this paper, we assumed that users who contribute to the system by mining cryptocurrencies generate negative externalities towards their mining peers. Indeed, the competition among miners is a very fundamental aspect of the mining process [1]. Nonetheless, by incorporating more miners, the blockchain becomes more robust [3]. Such robustness, in turn, may translate into an increase in the real value of the cryptocurrency under consideration [15, 13]. Therefore, by increasing the pool of miners, each miner is also contributing with positive externalities towards the system, and we leave such aspect as subject for future work.

Mining pools. Mining pools play a key role in today's public blockchain systems [2].¹ The competition analyzed in this paper applies to mining pools under two scenarios. First, from the perspective of the mining pool, it can use cloud resources for mining purposes. Therefore, the mining pools assume the role of players as considered in this work. Alternatively, the players are the end users, who contract mining pool services. Then, mining pools assume the role of ESPs. In the first case, we consider competition among mining pools, at the macro level, and in the latter case, we consider the micro-competition among end-users.

Multi-cryptocurrency ecosystem. In the cryptocurrency ecosystem, large mining pools typically decide, dynamically, which blockchain to mine. Such decisions are made based on different thresholds related to the value of the cryptocurrencies and the costs for mining (mining complexity). The churn of computational power across blockchains is a well-known source of price volatility, and different mechanisms have been developed to counteract migrations of miners across platforms [17]. One of those mechanisms is referred to as emergency difficulty adjustment (EDA), which reduces the difficulty of the puzzle when there are not many miners in the system, preventing the blockchain from dying.

Puzzle complexity. In this work, we assume that the time to mine the next block is monotonically decreasing with

¹For instance, <https://miningpoolhub.com/>.

respect to the number of miners. In Bitcoin, puzzle difficulty (complexity) is dynamically adjusted so that the time to mine a block varies between certain pre-established time bounds. Bitcoin target block generation rate is of 10 minutes.

In theory, due to the dynamic adjustment of puzzle complexity, Bitcoin throughput (number of blocks generated per time unit) does not depend on the number of miners. Nonetheless, in practice the throughput does vary over time and an increase in the number of miners can decrease the time between generation of blocks [9, 18]. In [5], the authors argue in favor of adjusting the frequency at which blocks are generated as a function of the congestion in the network. Our models can be extended to account for fixed and dynamic block generation rates, and we leave a further investigation of that subject for future work.

6. RELATED WORK

There is a vast literature investigating game theoretical aspects of blockchain systems [7, 16, 5]. Nonetheless, the literature on congestion games applied to such systems is scarce. In particular, to the best of our knowledge, there is no prior work investigating the competition at the network edge among miners as a congestion game, and its connection to multi-cryptocurrency markets.

Congestion games have been applied in the field of networking to account for security aspects [8], link congestion [6] and pricing of infrastructures and users [4]. In [5], the authors study Bitcoin as a congestion game, where the congestion occurs due to an increase in the number of transaction requests from users. In particular, the authors abstract away from several aspects of the competition between miners. In this paper, in contrast, we focus on the competition between miners.

Spiegelman et al. [16] adopted the framework of congestion games to model competition between miners of multiple cryptocurrencies who try to maximize utilities by choosing which puzzle (cryptocurrency) to mine [16]. The authors prove that there is no standard potential function for the game they propose, but that an ordinal potential always exists, implying that best response converges to a pure Nash equilibrium. Our work captures different aspects of the problem, and is complementary to [16]. An important similarity between the two works consists of establishing conditions under which pure Nash equilibria exist even when the game does not admit a standard potential function. The major differences between our work and [16] are: 1) in the modeling of the probability to succeed in solving a puzzle (see Section 5); 2) in the ESP decision, which is out of the scope of [16]; 3) in the action space (mining power), which is continuous in [16], precluding the use of crowding game results, and discrete in this paper, allowing us to rely on [10] to prove existence of pure Nash equilibria. We refer the reader to [16, 17] for additional references on the multi-cryptocurrency ecosystem and its security challenges.

7. CONCLUSION

We modeled the competition over several ESPs and over several blockchains characterizing multiple cryptocurrencies as a non-cooperative game. Then, we specialized our game to two cases: the ESP connection game and the cryptocurrency selection game. For each game, we showed properties

of the Nash equilibrium. In particular, leveraging results about congestion games, we establish the existence of pure Nash equilibria and characterize such equilibria through problems that admit efficient algorithmic solutions.

Acknowledgment Research conducted within the context of the THANES Associate Team, jointly supported by Inria (France) and FAPERJ (Brazil).

8. REFERENCES

- [1] DIMITRI, N. Bitcoin mining as a contest. *Ledger 2* (2017), 31–37.
- [2] EYAL, I. The miner’s dilemma. In *Security and Privacy* (2015), IEEE, pp. 89–103.
- [3] GARAY, J., KIAYIAS, A., AND LEONARDOS, N. The bitcoin backbone protocol. In *Theory and Applications of Crypto Techniques* (2015), Springer.
- [4] HASSIN, R., AND HAVIV, M. Equilibrium threshold strategies: The case of queues with priorities. *Operations Research 45*, 6 (1997), 966–973.
- [5] HUBERMAN, G., LESHNO, J. D., AND MOALLEMI, C. C. Monopoly without a monopolist. *SSRN* (2017).
- [6] JOHARI, R., AND TSITSIKLIS, J. N. Network resource allocation and a congestion game. In *Allerton* (2003).
- [7] KIAYIAS, A., KOUTSOPIAS, E., KYROPOULOU, M., AND TSELEKOUNIS, Y. Blockchain mining games. In *Conf. Economics and Computation* (2016), ACM.
- [8] MAILLÉ, P., REICHL, P., AND TUFFIN, B. Interplay between security providers, consumers, and attackers: a weighted congestion game approach. In *Conf. Decision Game Theory for Security* (2011), Springer.
- [9] MESHKOV, D., CHEPURNOY, A., AND JANSEN, M. Revisiting difficulty control for blockchain systems. In *Data Privacy Management, Cryptocurrencies and Blockchain Technology*. Springer, 2017, pp. 429–436.
- [10] MILCHTAICH, I. Congestion games with player-specific payoff functions. *Games and economic behavior 13*, 1 (1996), 111–124.
- [11] MILCHTAICH, I. Crowding games are sequentially solvable. *Intl. Journal Game Theory 27*, 4 (1998), 501.
- [12] MONDERER, D., AND SHAPLEY, L. S. Potential games. *Games and economic behavior 14*, 1 (1996), 124–143.
- [13] RAVAL, S. Ethereum Price Prediction, 2018. <https://tinyurl.com/ethereumpred>.
- [14] ROSENTHAL, R. W. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory 2*, 1 (1973), 65–67.
- [15] SHAH, D., AND ZHANG, K. Bayesian regression and Bitcoin. In *Allerton* (2014), IEEE, pp. 409–414.
- [16] SPIEGELMAN, A., KEIDAR, I., AND TENNENHOLTZ, M. Game of coins. *arXiv:1805.08979* (2018).
- [17] ULRICH, F. Attacking Bitcoin, 2017. <https://tinyurl.com/fulrich>.
- [18] WISDOM, B. Bitcoin Difficulty, 2018. <https://bitcoinwisdom.com/bitcoin/difficulty>.
- [19] XIONG, Z., FENG, S., NIYATO, D., AND WANG, P. Edge computing resource management and pricing for mobile blockchain. *arXiv:1710.01567* (2017).
- [20] ZHANG, Y., LIU, L., GU, Y., NIYATO, D., PAN, M., AND HAN, Z. Offloading in software defined network at edge with information asymmetry. *Journal of Signal Processing Systems 83*, 2 (2016), 241–253.