Using ‘Proof-of-Presence’ to Coordinate[*]

Adam Brandenburger[†]
Kai Steverson[‡]

This version: February 10, 2020

Abstract

Coordination in a distributed environment is a challenging problem. Here, we show how the idea of proof-of-work from computer science can be adapted to a game-theoretic setting to help players determine how many other players are present in a game and, therefore, available for coordination. The key idea is that a computational puzzle is distributed to all agents in the system and, if the puzzle is solved rapidly, then each agent can infer that, with high probability, a large number of other agents must have been working on it. We call this “proof-of-presence” (POP) and show how appending a POP mechanism to a game enables coordination when it would not otherwise happen, even if there are malicious agents present.

1 Introduction

Coordination among agents is a fundamental challenge in a distributed environment. In game theory, this issue is formalized as a coordination game among players who are spatially separated. The challenge is for the players to attain a Nash equilibrium of the game or – if there are Pareto-ranked equilibria – to attain an efficient equilibrium. In computer science, the issue of distributed coordination has been formalized as the Two Generals’ Problem (Gray, 1977) and the Byzantine Generals’ Problem (Lamport, Shostak, and Pease, 1982). In that problem, there are a number of generals who can safely attack a city only

[*] We are grateful to Christian Grewell, Hanna Halaburda, Margaret Minsky, Barry Nalebuff, Oliver Steele, and Marilyn Tsaih for valuable input, and to NYU Stern School of Business, NYU Shanghai, and J.P. Valles for financial support.

[†] Stern School of Business, Tandon School of Engineering, NYU Shanghai, New York University, New York, NY 10012, U.S.A., adam.brandenburger@stern.nyu.edu, http://www.adambrandenburger.com

[‡] DCI Solutions, Aberdeen, MD 21005, U.S.A., ksteverson@gmail.com

if sufficiently many other generals attack at the same time. Typically, it is assumed that communication among the generals is unreliable due to the presence of malicious agents who intercept messages or forge false ones. With the advent of blockchain (Nakamoto, 2008a), the idea of proof-of-work (POW) is often seen as a way to obtain coordination in these settings (Nakamoto, 2008b).[1]

This paper investigates another aspect of coordination, namely, uncertainty over the number of players who are present. In the language of the Byzantine Generals' Problem, each general might be confident the attack time is 4:00am (say this is the customary time for an attack) but less confident about how many other generals will be present to attack. (Other generals might have to prioritize other offensives.) In particular, we are interested in an anonymous environment, where there is no centralized list of players and no way to verify an individual's identity. For example, we can think of organizing a meeting using an anonymous message board. Anyone can post to the message board, and messages cannot be altered. However, there is no way to identify a post's author or how many people are posting. In this setting, a single malicious actor could easily pretend to be many people, artificially inflating how many attendees there promise to be. This effect thwarts any straightforward attempt to make a simple headcount.

In this paper, we show the POW mechanism can be adapted to make a reliable (albeit probabilistic) headcount.[2] We call this proof-of-presence (POP), and we analyze how it enables coordination. We construct a game-theoretic setting with an uncertain number of players, and analyze when a coordination equilibrium exists. We show that the POP mechanism enables coordination in settings where coordination would otherwise be impossible. Our main result establishes exact bounds on when coordination can be created through use of the POP mechanism.

Our POP mechanism works much like the original POW concept. A puzzle is posted on the message board that can only be solved by guesswork, but the solution can be immediately verified by any agent. Players who are present work on solving the problem, and post the solution if they find it. If the solution is found before a predetermined time, then players choose to take the coordination action (e.g., show up to the meeting); otherwise, they do not. The intuition for why this procedure provides POP is this: *If a puzzle can be solved only by guessing, and the rate of guessing is bounded (across agents), then a solution within a short period of time is a (probabilistic) proposition that a large number of agents must have been working on the puzzle.*

[1] POW has been explained as follows (Jakobsson and Juels, 1999): "[I]n a POW, a prover demonstrates to a verifier that she has performed a certain amount of computational work in a specified interval of time."

[2] Our procedure is inspired by the clever argument in DeDeo (2017), which describes the idea of POW via a parable of guests in a hotel who communicate with one another through an unreliable concierge.

We allow for both good-faith players (who participate in the game if a solution is found in time) and malicious agents (who do not participate in the game and may confuse the headcount). The intuition for the POP mechanism would seem to require the number of malicious agents (which may be probabilistic) to be small relative to the number of players. The argument would be that a small number of malicious agents could have only have a limited impact on when (or if) a solution is found. Hence, the puzzle provides a way of taking a (probabilistic) headcount that is immune to manipulation from sufficiently few such agents. Surprisingly, however, we are able to show that the POP mechanism functions even when the number of malicious agents is large, as long as the distribution of malicious agents satisfies a discrete form of log-concavity (a property satisfied by many common distributions). This property allows us to use monotone likelihood ratio property arguments to establish the effectiveness of POP in enabling coordination, regardless of the relative number of good-faith players and malicious agents.

2 The "Proof of Presence" Procedure

There is a universe consisting of good-faith players and malicious agents. There is a maximum of m good-faith players and a probability distribution $p = (p_1, p_2, \ldots, p_m)$ on the number of good-faith players.[3] There is a maximum of n malicious agents and a probability distribution $q = (q_0, q_1, \ldots, q_n)$ on the number of malicious agents. We will assume these distributions are full support over their ranges, independent, and that q follows a discrete form of log-concavity where $q_{j+1}/q_j < q_j/q_{j-1}$ for all $1 \leq j \leq n-1$. (We show how to weaken these assumptions later.) Both p and q are from the perspective of a good-faith player who is present, which is why 0 is not in the support of p. From now on, we will use the unqualified term "player" to refer to a good-faith player.

Each player has a coordination action and a null action (payoff W^*). The payoff from the coordination action is a strictly increasing function $W : \{1, 2, \ldots, m\} \to \Re$, where $W(k)$ is the payoff if k players take the coordination action. Malicious agents do not participate in the game (or can be thought of as always taking the null action). We will assume that $W(1) < W^*$ so that players do not want to take the coordination action by themselves.

Time is continuous and indexed by t in $[0, \infty)$. At time 0, a computational puzzle is posted to the anonymous message board and read by all players and malicious agents present. Each

[3] Harsanyi (1967-8) argued (though he did not formally prove) that uncertainty over whether or not a player is present in a game could be shifted to uncertainty over that player's payoff function (where the player is definitely present). The idea is to give the player a payoff of $-\infty$ from all actions (except a null action) to correspond to the case where s/he is absent. There seems no benefit (and an interpretational disadvantage) to our trying to follow this route.

player or malicious agent has a machine that works on the puzzle by randomly guessing, with a Poisson arrival rate λ of finding a solution. If a machine finds a solution, it posts it on the message board without notifying the solver. There is a pre-specified time T at which all players can read off the message board whether or not at least one solution has been found. Call this procedure the POP mechanism appended to the underlying game.[4]

We are interested in analyzing the strategy where a player who is present takes the coordination action if a solution is displayed at the pre-specified time T. Call this the coordination strategy. We want to know when this strategy is optimal in the sense of constituting a Nash equilibrium. In other words, we investigate when this strategy is optimal for each player assuming all other players follow it.

3 Probability Calculations

We now calculate probabilities, from the perspective of a player who is present. The probability that a total of k players and malicious agents are present, for $k = 1, 2, \ldots, m+n$, is given by

$$r_k = \sum_{i=\max\{1,k-n\}}^{\min\{k,m\}} p_i q_{k-i}.$$

By the properties of the exponential distribution, the probability that no solution is found by time T, if k players and malicious agents are working (independently) on the puzzle is $e^{-\lambda k T}$. It follows that the probability that exactly k players and malicious agents were working on the puzzle, if a solution is found by T, is given by

$$\phi_T(k) = \frac{r_k(1 - e^{-\lambda k T})}{\sum_{j=1}^{m+n} r_j(1 - e^{-\lambda j T})}.$$

Next, the probability that i players are present given k players and malicious agents are present is given by

$$\gamma_k(i) = \begin{cases} \frac{p_i q_{k-i}}{r_k} & \text{if } \max\{1, k-n\} \leq i \leq \min\{k, m\}, \\ 0 & \text{otherwise.} \end{cases}$$

Assume a solution is displayed at time T. Think of i and k as random variables repre-

[4] Note that we assume solutions are transmitted without the solvers' being aware they found a solution. This assumption preserves symmetry between the player(s) who solved and the other players, which is not necessary for any of our results but helps simplify the analysis. We could also allow players to see immediately that a solution has been posted. This would complicate our analysis but would not, we believe, lead to qualitatively different results.

senting the number of players and the number of players plus malicious agents, respectively. The expected payoff $\mathbf{E}_T W$ to a player following the coordination strategy after seeing the solution, assuming all other players do so as well, can be written as

$$\mathbf{E}_T W = \sum_{k=1}^{m+n} \phi_T(k) \sum_{i=1}^{m+n} \gamma_k(i) W(i).$$

By calculating limits as T approaches 0 (this uses l'Hôpital's Rule) and ∞, we find

$$\mathbf{E}_0 W \coloneqq \lim_{T\to 0} \mathbf{E}_T W = \frac{\mathbf{E}_{p,q}\left[W(i)k\right]}{\mathbf{E}_{p,q}\left[k\right]}$$

and

$$\mathbf{E}_\infty W \coloneqq \lim_{T\to\infty} \mathbf{E}_T W = \mathbf{E}_{p,q}\left[W(i)\right],$$

where $\mathbf{E}_{p,q}$ indicates the expectation taken under the a priori probability distributions p, q.

4 Results

We first present and discuss our results. We defer the proofs to the end of the section. The main mathematical result is the following.

Proposition 1. $\mathbf{E}_T W$ *is strictly decreasing in* T.

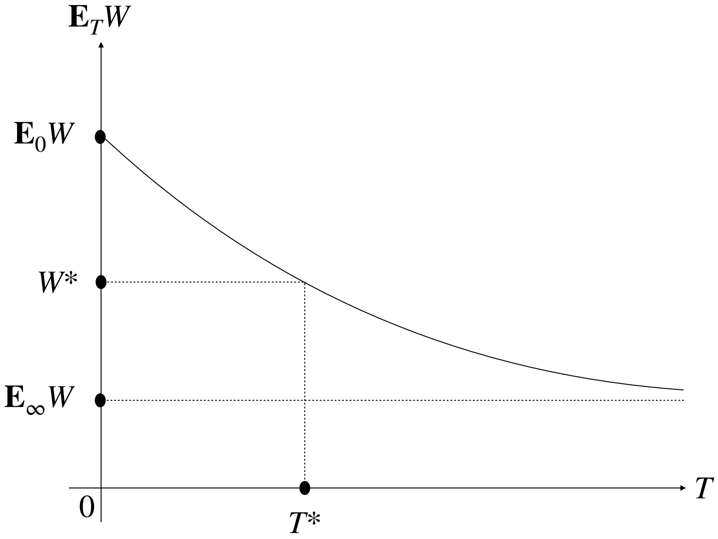

Figure 1

This result says that a player's payoff from following the coordination strategy strictly decreases in the pre-specified time T (when other players present follow this strategy). This validates the earlier intuition that a rapid solution to the computational puzzle should make choosing the coordination action more attractive to players. We can depict the general shape of $\mathbf{E}_T W$ as a decreasing function starting from $\mathbf{E}_0 W$ and decreasing to $\mathbf{E}_\infty W$, as shown in Figure 1 above.

Figure 1 assumes the case that $\mathbf{E}_\infty W < W^* < \mathbf{E}_0 W$, which is necessary to make the problem interesting. If $W^* > \mathbf{E}_0 W$, then even a puzzle that was immediately solved could not induce the players to take the coordination action. This would mean coordination is impossible. Conversely, if $W^* \leq \mathbf{E}_\infty W$, then the puzzle is not necessary for coordination to be an equilibrium.

Figure 1 illustrates how Proposition 1 can be used to determine the conditions under which the coordination strategy is a Nash equilibrium. Let $0 < T^* < \infty$ solve $\mathbf{E}_T W = W^*$. Fix a pre-specified time T and a particular player i. Assume all other players present are playing the coordination strategy (with respect to T). If a solution is revealed at T, then the expected payoff of playing the coordination strategy to player i is $\mathbf{E}_T W$. It will therefore be optimal for the player to take the coordination action exactly when $\mathbf{E}_T W \geq \mathbf{E}_{T^*} W$, which, by Proposition 1, occurs when $T \leq T^*$. If no solution is revealed at T, then player i will take the null action since we assumed $W(1) < W^*$. Therefore, $T \leq T^*$ is the precise condition that makes the coordination strategies form a Nash equilibrium

We now state this more formally.

Proposition 2. *Suppose $\mathbf{E}_\infty W < W^* < \mathbf{E}_0 W$. Let $0 < T^* < \infty$ solve $\mathbf{E}_T W = W^*$. Append the POP mechanism with a pre-specified time T, and consider the strategy profile where all players choose the coordination action if a solution is revealed at time T, and the null action otherwise. Then it is a Nash equilibrium for all players to choose this strategy if and only if $T \leq T^*$. Moreover, without the POP mechanism the only Nash equilibrium is where all players always take the null action.*

The first part of Proposition 2 follows from Proposition 1 using the logic we outlined above. To see the second part, note that, without any mechanism, $\mathbf{E}_\infty W$ is the payoff to a player from taking the coordination action assuming all other present players take this action, too. This represent the most optimistic assumption about the strategies of the other players, and thus $\mathbf{E}_\infty W$ is an upper bound on a player's expected payoff from taking the coordination action. But Proposition 2 assumes that $W^* > \mathbf{E}_\infty W$, which implies that, without a mechanism, the null action is always better.

We now turn to proving Proposition 1. Given two distributions p_1, p_2, we will use the notation $p_1 \succ_D p_2$ to indicate that p_1 first-order stochastically dominates (FOSD) p_2. To prove Proposition 1 we will prove two statements: (1) $\gamma_{k+1} \succ_D \gamma_k$ for all k; and (2) $\phi_{T+1} \prec_D \phi_T$ for all T. In words, statement (1) says that a player's beliefs about the number of players present increases with the total number of players plus malicious agents. Statement (2) says that a player's beliefs about the total number of players plus malicious agents decreases with T. Intuitively, these statements combine to deliver Proposition 1. To see this formally, define

$$V(k) := \sum_{i=1}^{m+n} \gamma_k(i) W(i),$$

so that we can rewrite $\mathbf{E}_T W$ as

$$\mathbf{E}_T W = \sum_{k=1}^{m+n} \phi_T(k) V(k).$$

Statement (1), along with the fact that W is strictly increasing in i, implies that V is strictly increasing in k. Statement (2), along with the fact that V is strictly increasing in k, implies that $\mathbf{E}_T W$ is strictly decreasing in T. Thus, to prove Proposition 1, it suffices to prove statements (1) and (2).

To do so, we will prove the stronger statements: (3) γ_k has the increasing monotone likelihood ratio property (MLRP) relative to k; and (4) ϕ_T has the decreasing MLRP relative to T. Using the standard fact that the increasing (decreasing) MLRPimplies increasing (decreasing) first-order stochastic dominance, we see that (3) implies (1), and (4) implies (2).

Lemma 1. *γ_k satisfies the increasing MLRP with respect to k.*

Proof. The definition of MLRP requires that for all i and $1 \leq k \leq m+n-1$

$$\frac{\gamma_{k+1}(i)}{\gamma_k(i)} \leq \frac{\gamma_{k+1}(i+1)}{\gamma_k\left(i+1\right)}, \tag{1}$$

with the inequality holding strictly for at least one i. (This requirement that Equation (1) holds strictly at least once is not always included in the definition of MLRP, but it is necessary to get that MLRP implies FOSD.)

To avoid issues of dividing by 0, we instead we write this as

$$\gamma_{k+1}(i)\gamma_k(i+1) \leq \gamma_k(i)\gamma_{k+1}(i+1).$$

If $i \geq \min\{k, m\}$, then $\gamma_k\left(i+1\right) = 0$, and the right-hand side is always non-negative, hence the inequality holds.

Next suppose i is in the range $[\max\{1, k-n\}, \max\{1, k+1-n\})$. Here $\gamma_{k+1}(i) = 0$ which again ensures the inequality always holds.

Now consider i in the range $[\max\{1, k+1-n\}, \min\{k, m\})$. We no longer have to worry about dividing by 0 so we revert to considering our original inequality (1). Here,

$$\frac{\gamma_{k+1}(i)}{\gamma_k(i)} = \frac{p_i q_{k+1-i}/r_{k+1}}{p_i q_{k-i}/r_k} = \frac{q_{k+1-i}}{q_{k-i}} \times \frac{r_k}{r_{k+1}}.$$

By our assumption on q, we conclude that q_{k+1-i}/q_{k-i} is strictly increasing in i. It follows that for i in the complete range $[\max\{1, k-n\}, \min\{k, m\})$, the desired inequality holds strictly. This establishes that γ_k satisfies the increasing MLRP.

We now need to show that MLRP holds strictly for at least one i. If $[\max\{1, k-n\}, \max\{1, k+1-n\})$ contains a whole number i, then $\gamma_k(i) \neq 0$ by the full-support assumption and $\gamma_{k+1}(i) = 0$ as discussed above, which gives the desired strict inequality. On the other hand, if $[\max\{1, k-n\}, \max\{1, k+1-n\})$ does not contain a whole number i, then it must be that $[\max\{1, k+1-n\}, \min\{k, m\})$ contains a whole number i. And for this i we showed above the MLRP holds strictly. □

Lemma 1 is the only place where we use the assumption that p and q are full support and independent, or the assumption that q is (discrete) log-concave. We could do away with both assumptions by instead directly assuming that γ_k is increasing, in the sense of FOSD, in k. In words, this assumption would mean that the estimate of the number of players is increasing in the total number of players plus malicious agents.

Lemma 2. *$\phi_T(k)$ satisfies the decreasing MLRP in T and $\phi_T \neq \phi_{T'}$ for $T \neq T'$.*

Proof. Assume $T' \geq T > 0$, and we want to show that $k = 1, 2, \ldots, m+n-1$,

$$\frac{\phi_{T'}(k+1)}{\phi_T(k+1)} \leq \frac{\phi_{T'}(k)}{\phi_T(k)},$$

or equivalently that

$$\frac{\phi_{T'}(k+1)}{\phi_{T'}(k)} \leq \frac{\phi_T(k+1)}{\phi_T(k)},$$

with the inequality hold strictly for at least one k. Accordingly, we will show that $\phi_T(k+1)/\phi_T(k)$ is strictly decreasing in T. Let

$$g(T, k, \lambda) = \frac{1 - e^{-\lambda(k+1)T}}{1 - e^{-\lambda k T}}.$$

Then, by the definition of $\phi_T(k)$,

$$g(T,k,\lambda) \times \frac{r_{k+1}}{r_k} = \frac{\phi_T(k+1)}{\phi_T(k)},$$

so it is enough to show that $g\left(T,k,\lambda\right)$ is strictly decreasing in T. We can calculate

$$\frac{\partial g\left(T,k,\lambda\right)}{\partial T} = \frac{\lambda(k+1)e^{-\lambda(k+1)T}}{1-e^{-\lambda kT}} - \frac{\lambda k e^{-\lambda kT}(1-e^{-\lambda(k+1)T})}{(1-e^{-\lambda kT})^2},$$

from which we find $\partial g/\partial T < 0$ if and only if

$$h\left(T,k,\lambda\right) := (k+1)e^{-\lambda T} - e^{-\lambda(k+1)T} - k < 0.$$

Since $h(T,0,\lambda) = 0$, it is enough to show that $\partial h(T,k,\lambda)/\partial k < 0$ for all $k > 0$. Calculating

$$\frac{\partial h\left(T,k,\lambda\right)}{\partial k} = e^{-\lambda T} + \lambda T e^{-\lambda(k+1)T} - 1.$$

Suppose, by way of contradiction, that $\partial h/\partial k \geq 0$. Then

$$\frac{e^{\lambda(k+1)T} - e^{\lambda kT}}{\lambda T} \leq 1.$$

By the mean value theorem, there is then a number c, where $\lambda kT \leq c \leq \lambda(k+1)T$, such that $e^c \leq 1$, which is impossible since $c > 0$. □

5 Discussion

In this paper, we showed the POW concept can be adapted to provide a reliable, albeit probabilistic, headcount in an anonymous environment with malicious agents. We refer to this adaptation as the proof-of-presence (POP) mechanism. The key idea is that if a puzzle can only be solved through guessing, then the faster it is solved the more likely it is that more people are working on it. The primary application we have in mind is organizing meetings for anonymous online communities in the presence of potential disruptors who attempt to make any headcount as unreliable as possible. In this environment, our analysis shows that the POP mechanism provides a tool that can help the community members successfully coordinate on a decision to meet.

For simplicity, we ignored a number of complicating factors that may appear in the real world. We now discuss some such factors and how they would impact our analysis of the POP mechanism. The first factor is that players or malicious agents may apply different levels of computational power to guessing the puzzle, leading to different arrival rates for a solution. This seem especially likely in the case of the malicious agents who, depending on the cost, have an incentive to apply additional computational power to disrupt coordination. However, we can simply think of each malicious agent in our setting as representing a unit of computational power applied by the malicious agents. Under this re-conceptualization, all of our results remain the same and the POP mechanism still functions. This re-conceptualization works because the malicious agents only matter through the impact they have on the solution time for the puzzle. Whether one or many malicious agents are involved is irrelevant, assuming they apply the same aggregate computational power. For the players, since they want the POP mechanism to succeed, they can simply all agree to use the same level of computational power. Alternatively, we could vary the arrival rates (the λ's) by player to model different levels of computational power. We conjecture that this would only serve to complicate the analysis without qualitatively changing the results.

The reinterpretation of malicious agents as units of computational power highlights the surprising feature that the POP mechanism does not depend on there being only a small number of malicious agents present. This means that even if the malicious agents have far more computational power than the players have, the POP mechanism still functions. This is because behavior of the malicious agents does not interfere with the inference that a faster solution means there are likely more players present. Increasing the (probabilistic) amount of malicious computational power only lowers the threshold on how fast the solution must be solved to make coordination worthwhile. It does not change the fact that such a threshold exists.

A second complicating factor not addressed in our analysis is that solving the puzzle may be costly for the players. This gives an incentive for players to free-ride by not spending any effort in solving the puzzle while still joining in the coordination game. We can solve this issue the same way it is solved in POW mechanisms such as bitcoin, by providing a reward to whomever posts the first solution. To maintain anonymity, the reward could be in the form of a gift card messaged to the winner's anonymous account.

Another factor our analysis does not address is that of the optimal value of T. The optimal T must balance three considerations. First, the value of T must be below the threshold T^* to ensure that the coordination strategies constitute an equilibrium, as characterized in the results of this paper. Second, the mechanism wants to minimize the number of times players take the coordination action when there are not enough players present to justify this choice. Third, the mechanism wants to maximize the number of times the players take the coordination action when there are enough players present to justify this choice. The second consideration pushes in favor of lowering the cutoff T, while the third consideration pushes in favor of raising T. Finding the balance of these considerations looks to be a fairly standard mechanism design problem where the optimal T will depend on the payoffs $W(\cdot)$ and W^*. Also, there may be additional payoff features not included in our analysis. For example, a successful coordination outcome may have positive externalities that accrue to the non-present players. Characterizing the optimal T is beyond the scope of our current paper and is left for future work.

References

DeDeo, S., "The Bitcoin Paradox: Why Cryptocurrency Will Always Be Political," 2017, at `http://nautil.us/issue/55/trust/the-bitcoin-paradox`.

Gray J., "Notes on Database Operating Systems," IBM Research Report RJ 2188, 1978.

Harsanyi, J.,"Games with Incomplete Information Played by 'Bayesian' Players, I-III," *Management Science*, 14, 1967-8, 159-182, 320-334, 486-502.

Jakobsson, M., and A. Juels, "Proofs of Work and Bread Pudding Protocols," in *Secure Information Networks: Communications and Multimedia Security*, Kluwer Academic Publishers, 1999, 258-272.

Lamport, L., R. Shostak, and M. Pease, "The Byzantine Generals Problem," *ACM Transactions on Programming Languages and Systems*, 4, 1982, 382-401.

Nakamoto, S., "Bitcoin: A Peer-to-Peer Electronic Cash System," 2008a, at https://bitcoin.org/bitcoin.pdf.

Nakomoto, S., "The Proof-of-Work Chain is a Solution to the Byzantine Generals' Problem," 2008b, at https://satoshi.nakamotoinstitute.org/emails/cryptography/11/.