

# Full Analysis of Nakamoto Consensus in Bounded Delay Networks

Juan Garay<sup>1</sup>, Aggelos Kiayias<sup>2</sup>, and Nikos Leonardos<sup>3</sup>

<sup>1</sup> Texas A&M University, [garay@tamu.edu](mailto:garay@tamu.edu)

<sup>2</sup> University of Edinburgh & IOHK, [akiayias@inf.ed.ac.uk](mailto:akiayias@inf.ed.ac.uk)

<sup>3</sup> National and Kapodistrian University of Athens, [nikos.leonardos@gmail.com](mailto:nikos.leonardos@gmail.com)

**Abstract.** Nakamoto consensus, arguably the most exciting development in distributed computing in the last few years, is in a sense a recasting of the traditional state-machine-replication problem in an unauthenticated setting, where furthermore parties come and go without warning. The protocol relies on a cryptographic primitive known as *proof of work* (PoW) which is used to throttle message passing with the PoW difficulty level being appropriately adjusted throughout the course of the protocol execution.

While the original formulation was only accompanied by rudimentary analysis, significant and steady progress has been made in abstracting the protocol's properties and providing a formal analysis under various restrictions and protocol simplifications. Still, a full analysis of the protocol that includes its target recalculation and timestamp adjustment mechanisms which equip it to operate in its intended setting of bounded communication delays, imperfect clocks and dynamic participation, has remained open.

This paper's main result fills this gap presenting a proof that Nakamoto's protocol achieves, under suitable conditions, consistency and liveness in the above setting. A number of technical tools and techniques are introduced that may be of independent interest in the analysis of blockchain protocols.

## 1 Introduction

Nakamoto's blockchain protocol [23] is a consensus protocol where parties engage in the collection and organization of transactions in a ledger without having any information about each other or even precise knowledge of the number of parties running the protocol at any given time. This is in contrast to classical models and results in consensus (aka Byzantine agreement) [25, 19] and other fundamental distributed computing tasks, where it is typically assumed that parties have pairwise authenticated communication channels or are initialized with the public keys of all participants. Instead, Nakamoto's blockchain protocol relies on the cryptographic primitive known as *proof-of-work* (PoW, aka cryptographic puzzles) [10, 27, 2, 15]), to throttle message transmission and stochastically create opportunities for unifying the parties' possibly diverging views, despite the presence of a subset of them acting adversarially.

Given that the original protocol was presented with only a rudimentary analysis focusing solely on the application context of fund transfers, a number of works have attempted to isolate the protocol’s properties and provide a formal analysis. The first analysis, presented in [12], focused on a synchronous execution model, and assuming the probability of the parties to solve a PoW over a single message-passing round is suitably restricted, proved that the protocol satisfies consistency and liveness as long as the total computational power in the system is in favor of the honest parties. Two limitations of this first analysis were that the target recalculation mechanism of the blockchain protocol which adjusts the hardness of PoWs was excluded and that the execution model considered synchronous communication rounds.

Addressing the latter problem was undertaken in [24] (with further improvements in follow-up works [18, 26]), where the blockchain protocol was analyzed in the so-called *bounded-delay model* (cf. [9])<sup>4</sup>, showing the protocol secure for a favorable choice of network delay  $\Delta$  with respect to its hard-coded PoW hardness parameter, and its insecurity in the general case where  $\Delta$  is chosen adversarially. Technically, the main challenge to address in transitioning to the bounded-delay model is the fact that the usefulness of a certain event in the protocol execution (e.g., the creation of a PoW at time  $t$ ) is affected by events that are happening at times up to  $t + \Delta$  forward in time (e.g., the creation of another PoW) and hence this dependence asks for additional care to be applied in the probabilistic analysis.

The problem of analyzing the target recalculation mechanism was addressed in [13], albeit again in the synchronous communication model, by introducing a setting where parties’ participation is allowed to change round by round following a predetermined schedule that has a bounded rate of change. The main technical difficulty addressed in that setting was the fact that PoW successes are not independent events in the execution since the difficulty of the PoW primitive is determined by preceding execution events instead of being fixed throughout, as in [12, 24].

While the above works have significantly improved our understanding regarding the behavior of Nakamoto consensus in successively more refined theoretical models, the full analysis of the actual protocol has been elusive. Specifically, all previous works analyzed simplified versions of the protocol removing or adjusting protocol elements that deal with bounded delay networks and fluctuating participation. For instance, all previous works ignore the way the protocol adjusts local clocks based on on-chain timestamps and the actual timestamp rules for incoming blocks, mechanisms that Nakamoto (presumably) included to deal with the fact that no protocol can realistically assume perfectly synchronized clocks in an imperfect network. Manipulating timestamps and exploiting the

---

<sup>4</sup> In this model, formulated by Dwork, Lynch and Stockmeyer [9], there is an upper bound  $\Delta$ , unknown to the protocol, in the delay that the adversary may inflict on the delivery of messages.

protocol’s timestamp validation mechanisms are a well known attack vector in the Bitcoin community.<sup>5</sup>

Thus, the question is whether Nakamoto’s consensus—with all its adjustment mechanisms included—retains its properties in a bounded-delay network. Importantly, we want to answer this question when the parties’ number is dynamically changing without following a predetermined schedule, i.e., it is *adaptively* selected by the adversary, possibly even reacting to events that happen in the protocol execution, as long as the rate of change is bounded by a constant.

**Our results.** Our main result is the proof that Nakamoto’s protocol achieves consistency and liveness in bounded-delay networks with adaptive dynamic participation. While our results to some extent draw from the previous analyses of [12, 24, 13], technically our objective is significantly more challenging, and new techniques and insights are needed in order to realize it.

Importantly, our work for the first time takes into account the way that Nakamoto’s protocol adjusts local clocks using on-chain timestamps and validates incoming block timestamps. Specifically, the protocol allows incoming block timestamps in the near future, as determined by a protocol parameter and rejects blocks that have a timestamp in the past of the median time of a specific number of blocks on-chain—also a protocol parameter. These mechanisms open up an array of attack vectors that now should be considered in the security proof.

At the core of our analysis is the function  $f(T, n)$ , which determines the probability that  $n$  parties executing the protocol at a certain time find a PoW whose difficulty is determined by the “target”  $T$  (here, without loss of generality,  $n$  equates the number of parties with the number of CPUs of equal power running the protocol). Given that  $n$  is unknown and continuously changing, Nakamoto’s protocol adjusts  $T$  at regular intervals called *epochs*. As we show, the protocol’s resilience to attacks will stem from its ability to keep  $f(T, n)$  close to a suitable value that is favorably positioned with respect to the, otherwise unknown, network delay  $\Delta$ . Starting with the assumption that the protocol is initiated at an appropriate  $f$  value, the blockchain protocol will recalculate  $T$  to approximate that initial value by estimating the number of active parties  $n$  per epoch. The estimation is based on the observed production of PoWs as recorded in the blockchain itself and the relative timings of their production. A complication here is that timestamps may be manipulated in various ways during the protocol execution and this is something that the analysis should take into account.

The first major technical challenge is to work in a probabilistic setting where the random variables corresponding to the *cumulative difficulty* of PoWs (rather than their number) collected by the protocol participants capture the adaptive dynamic evolution of participants as well as the fact that some of these variables’ values at a certain round may be affected by events in the future. The latter issue asks for a lower bound estimation of the aggregate difficulty (in terms of PoWs

---

<sup>5</sup> E.g., timejacking and poison-pill attacks or difficulty-raising by timestamp manipulation, see [http://culubas.blogspot.com/2011/05/timejacking-bitcoin\\_802.html](http://culubas.blogspot.com/2011/05/timejacking-bitcoin_802.html), and <https://bitcointalk.org/index.php?topic=43692.msg521772#msg521772>.

of different targets) collected by the honest parties over a period of time that is “isolated” from any future PoW event for a period of  $\Delta$  rounds. At the same time, we need an upper bound on the aggregate difficulty amassed by adversarial parties while accounting for the fact that the adversary may choose to work on very difficult PoW instances for which it will be impossible to control their stochastic advantage via concentration bounds, due to high variance. We address this by introducing a suitable concept of “typical” execution where concentration bounds can be meaningfully applied to the relevant random variables.

The second major technical challenge is the analysis of the clock adjustment mechanism and validation of incoming block timestamps, and the rule for determining the timestamp for the next block (namely, that the new time stamp must be greater than the median of a certain number of blocks—see Section 3). The correctness of the rule critically depends on the existence on protocol executions of “winning streaks” by honest parties. Accordingly, we call such executions *hot hand* and we demonstrate how they result in regular high concentrations of honestly contributed blocks, that is necessary to ensure that honest parties’s chains have timestamps that move forward. Our analysis also reveals a property of the Bitcoin blockchain that was before not formally understood and can be seen as a strengthening of chain quality (CQ) [12]: honest chains are not only guaranteed to have regular contributions by honest parties as CQ dictates, but these contributions, with overwhelming probability, will come in clusters, so that regularly it will happen that in a sequence of  $2k_{\text{med}} - 1$  consecutive blocks at least  $k_{\text{med}}$  blocks will be contributed by the honest parties, where  $k_{\text{med}}$  is a small constant. Prior to our work such statements could only be proved for large values of  $k_{\text{med}}$  bounding the adversary by  $1/3$  and hence they were unsuitable for arguing the security of Nakamoto’s parameterization which imposes a lower bound on a block timestamp by the median of the last  $2k_{\text{med}} - 1 = 11$  blocks and seeks to argue protocol security for an adversarial bound below  $1/2$ .

Putting together the properties of typical and hot-hand executions we distill a set of conditions, stated here at a high level (refer to Section 4.2 for the detailed description), under which consistency and liveness of Nakamoto consensus can be shown. First we need honest majority, i.e., the honest parties’ hashing power exceeds the adversarial parties’ power. Second, we need epochs to be long enough to be able to adequately measure the change in the hashing power so it is reflected in the target recalculation (refer to C1 in Section 4.2). Then we need an upper bound on the network delay in terms of the block production rate and the other parameters (refer to C2 in Section 4.2) as well as a bound on the rate the parties dynamically change over time (C3 in Section 4.2).

*Remark.* The reader might wonder why the number of parties is allowed to adaptively dynamically evolve following even an exponential increase while the upper bound on the delay  $\Delta$  is assumed fixed throughout the execution and not allowed to dynamically evolve as well. The reason is that Bitcoin was designed with exactly this setting in mind. (To see evidence of this consider that PoW production in the protocol is fixed to be at 10 min intervals, irrespectively of the computational hashing power available to the network which has increased



significantly—and at periods of times even exponentially— since its initiation in 2009.) Designing a protocol that can absorb changes in  $\Delta$  as well is an interesting, albeit rather theoretical, open question.

Other directions for future work include proving the tightness of the necessary conditions for security. Recent work has achieved such tight bounds but for simplified variants of Nakamoto’s protocol—without target recalculation and timestamp adjustments (cf. [14, 7]). Finally, improving Nakamoto’s protocol so that it does not depend on an external clock and it becomes its own timekeeper would be interesting to demonstrate. Combining permissionless synchronization techniques developed in the proof-of-stake setting [3] with the 2-for-1 PoW technique of [12] would be a possible way forward.

**Differences with previous versions.** The current version of the paper significantly extends the analysis to include clock adjustment and blockchain timestamp validation mechanisms consistent with Nakamoto’s implementation.

**Organization of the paper.** The remainder of the paper is organized as follows. In Section 2 we present the network, protocol execution and adversarial model; in particular, we define the *dynamic bounded-delay setting* where our analysis is performed. In Section 3 we present Nakamoto’s consensus protocol, including its target recalculation and timestamp calculation functions, followed by the desired properties the blockchain protocol should satisfy. Section 4 contains the bulk of our analysis, in particular the relevant protocol parameters and the conditions mentioned above under which Nakamoto consensus can be achieved, followed by characterizations of executions that are typical (Section 4.4) and hot-hand (Section 4.5), where the basic blockchain properties are satisfied (4.7), from which proofs of consistency and liveness can be derived (Section 4.8).

Some standard mathematical facts, detailed protocol descriptions and some of the proofs are presented in the appendix.

## 2 Model

We describe our protocols in a “full” partially synchronous model where *both* communication and processors are partially synchronous. Specifically, in the model there is an upper bound  $\Delta$  in the delay (measured in number of rounds) that the adversary may inflict to the delivery of any message, and an upper bound  $\Phi$  on the potential difference in party’s clocks (also measured in number of rounds—see below). The precise values of  $\Delta$  and  $\Phi$  will be *unknown* to the protocol (and in particular regular protocol participation will not rely on using  $\Delta$  or  $\Phi$  as a time-out parameter). However, the security of the protocol will be dependent on how specific protocol parameters relate to  $\Delta$  and  $\Phi$  in ways we will explicitly define. Observe that “rounds” still exist in the model, but now these are not synchronization rounds where messages are supposed to be delivered to honest parties. Next, we adapt Canetti’s formulation of “real world” notion of protocol execution [4–6] for multi-party protocols, to the dynamic setting with a varying number of parties, bounded delays and partially synchronized clocks.

**Round structure and protocol execution.** As in [12, 24], the protocol execution proceeds in “rounds” with inputs provided by an environment program denoted by  $\mathcal{Z}$  to parties that execute the protocol  $\Pi$ . The adversary  $\mathcal{A}$  is “adaptive,” and allowed to take control of parties on the fly, as well as “rushing,” meaning that in any given round the adversary gets to observe honest parties’ actions before deciding how to react. Network and hash function access is captured by a diffusion (“gossiping”) functionality and a random oracle (RO) functionality, respectively (see below). The diffusion functionality is similar to those in [12, 24]; it allows the order of messages to be controlled by  $\mathcal{A}$  (i.e., there is no atomicity guarantees when multiple messages are sent) and, furthermore, the adversary is allowed to “spoof” the source information on every message (i.e., communication is not authenticated).  $\mathcal{A}$  can inject messages for selective delivery but cannot change the contents of the honest parties’ messages nor prevent them from being delivered beyond  $\Delta$  rounds of delay—a functionality parameter.

The environment program  $\mathcal{Z}$  determines the protocol execution; it creates and interacts with other instances of programs at the discretion of a control program  $C$ . Following [5],  $(\mathcal{Z}, C)$  forms of a *system of interactive Turing machines* (ITM’s). The only instances allowed by  $C$  are those of the protocol program  $\Pi$ , an adversary  $\mathcal{A}$ . These are called ITI’s (interactive Turing Machines Instances). We refer to [5] for further details on the mechanics of these aspects of the model. The only additional feature that is relevant to our setting is that we assume each instance is initialized with a special Boolean flag denoted by `active` which is set to false upon initialization.

**Functionalities available to the parties.** We next present the functionalities that are available to all parties running the protocol and the adversary and abstract the hash function (RO), the network and parties’ clocks. Note that the functionalities below share common state and realizing them by other protocols is outside the scope of the present work; in our exposition, they merely capture explicitly the assumptions we make about our execution model.

- *The RO functionality.* It accepts queries of the form (`Compute`,  $x$ ) and (`Verify`,  $x, y$ ). For the first type of query, assuming  $x$  was never queried before, a value  $y$  is sampled from  $\{0, 1\}^\kappa$  and it is entered to a table  $T_H$ . If  $x$  was queried before the pair  $(x, y)$  is recovered from  $T_H$ . The value  $y$  is provided as an answer. For the second type of query, a membership test is performed on the table. Honest parties are allowed to ask *one query per round* of the type `Compute` and unlimited queries of the type `Verify`.<sup>6</sup> The adversary  $\mathcal{A}$  is given a bounded number of `Compute` queries per round and no `Verify` queries (the adversary can easily simulate those locally). The bound for the adversary is determined as follows. Whenever a corrupted party is activated the bound is increased by 1; whenever a query is asked the bound is decreased by 1 (it is not necessary that the specific corrupted party makes the query).

---

<sup>6</sup> Note that we exclude denial-of-service attacks from our modeling, where  $\mathcal{A}$  depletes the running time of parties by sending them too many messages for verification.

- *The diffusion functionality.* Message passing and round bookkeeping is maintained by this functionality. A round variable *globalclock* is initialized to 0. For each party a string denoted by `RECEIVE()` is maintained and the party is allowed to fetch the contents of its corresponding `RECEIVE()` at any time. The functionality records all messages of the form `(Diffuse, m)` it receives from the parties. Completion of a round for a party is indicated by sending a special message `(RoundComplete)`. The adversary  $\mathcal{A}$  is allowed to receive all the currently recorded `Diffuse` messages at any time and messages to the `RECEIVE()` strings as desired. The round is completed when the adversary submits its `(RoundComplete)` message. In such case, the functionality inspects the contents of all `RECEIVE()` strings and includes any messages  $m$  that were diffused by the parties  $\Delta$  rounds ago but not contributed by the adversary to the `RECEIVE()` tapes (in this way guaranteeing message delivery up to  $\Delta$  rounds). It also flushes any diffuse records that are placed in the `RECEIVE()` string of all parties. The variable *globalclock* is then incremented and a new round begins.
- *The clock functionality.* This functionality, parameterized by  $\Phi$ , maintains parties clocks' values within this bound with respect to global time. The parties use `(RequestTime)` to obtain its local clock value from the functionality. The adversary  $\mathcal{A}$  is allowed to issue `(AddShift, P, d)`, with  $d \in [-\Phi, \Phi]$ , commands to the functionality, in order to alter party  $P$ 's clock value, so long as  $|d| \leq \Phi$ , i.e., honest parties' clocks are roughly synchronized.<sup>7</sup>

**The dynamic partially synchronous setting.** Given the functionalities as described above observe that contrary to prior formalizations, the adversary can choose the termination of the round thus deciding on the spot how many honest parties were activated adaptively. (In previous works [12, 24, 13] the adversary is restricted to a preset number of activations.) In each round, the number of parties that are active in the protocol is denoted by  $n_r$  and is equal to the total number of parties that have submitted the `(RoundComplete)` indicator to the diffusion functionality and have their internal flag `active` set to true. Determining  $n_r$  can only be done by examining the view of all honest parties and is not a quantity that is accessible to any of the honest parties individually. The number of parties controlled by  $\mathcal{A}$  in a round  $r$  is similarly denoted by  $t_r$ .

Parties, when activated, are able to read their input tape `INPUT()` and communication tape `RECEIVE()` from the diffusion functionality. If a party finds that its `active` flag is false, it enters a “bootstrapping” mode where it will diffuse a discovery message and synchronize with the rest of the active parties in the network (in the case of Nakamoto consensus, the party will send a request for the latest blockchains, will collect all of them until a time-out parameter is reached and then will pick the most difficult one to start mining).<sup>8</sup> When the synchronization phase terminates, the party will set its `active` flag to true and after

<sup>7</sup> Refer to [16] for further details on clock functionalities.

<sup>8</sup> Refer to Section 4.2 (specifically, the discussion on the  $\Delta_{\text{bootstr}}$  parameter) for an adequate time-out value depending on the epoch's length and probability of at least one honest party out of the initial number of parties solving a PoW.

this point it will be counted among the honest parties. An honest party goes “offline” when it misses a round, i.e., the adversary issues a (`RoundComplete`) but that party misses the opportunity to complete its computation. To record this action, whenever this happens we assume that the party’s `active` flag is set to false (in particular this means that a party is aware that it went offline; note, however, that the party does not need to report it to anyone). Also observe that parties are unaware of the set of active parties. As in previous works (e.g., [12]), we assume, without loss of generality, that each honest party has the same computational power.<sup>9</sup>

We will restrict the environment to fluctuate the number of parties in a certain limited fashion. Suppose  $\mathcal{Z}$  with fixed coins produces a sequence of parties  $n_r$  where  $r$  ranges over all rounds of the execution. We define the following property, which is a finite-sequence version of a similar property introduced in [13] for infinite sequences.

**Definition 1.** For  $\gamma \in \mathbb{R}^+$  we call  $(n_r)_{r \in [0, B)}$ , where  $B \in \mathbb{N}$ ,  $(\gamma, s)$ -respecting if for any set  $S \subseteq [0, B)$  of at most  $s$  consecutive integers,  $\max_{r \in S} n_r \leq \gamma \cdot \min_{r \in S} n_r$ .

We say that  $\mathcal{Z}$  is  $(\gamma, s)$ -respecting if for all  $\mathcal{A}$  and coins for  $\mathcal{Z}$  and  $\mathcal{A}$  the sequence of parties  $n_r$  is  $(\gamma, s)$ -respecting.

The term  $\{\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}^P(z)\}_{z \in \{0, 1\}^*}$  denotes the random variable ensemble describing the view of party  $P$  after the completion of an execution running protocol  $\Pi$  with environment  $\mathcal{Z}$  and adversary  $\mathcal{A}$ , on input  $z \in \{0, 1\}^*$ . We consider a “standalone” execution without any auxiliary information and we will thus restrict ourselves to executions with  $z = 1^\lambda$ . For this reason we will simply refer to the ensemble by  $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}^P$ . The concatenation of the view of all parties ever activated in the execution is denoted by  $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}$ .

**Properties of protocols.** In our theorems we will be concerned with *properties* of protocols  $\Pi$  running in the above setting. Such properties will be defined as predicates over the random variable  $\text{VIEW}_{\Pi, \mathcal{A}, \mathcal{Z}}$  by quantifying over all possible adversaries  $\mathcal{A}$  and environments  $\mathcal{Z}$ . Note that all our protocols will only satisfy properties with a small probability of error in  $\lambda$  as well as in a parameter  $k$  that is selected from  $\{1, \dots, \lambda\}$  (with foresight, we note that in practice one would be able to choose  $k$  to be much smaller than  $\lambda$ , e.g.,  $k = 6$ ).

### 3 The Nakamoto (Ledger) Consensus Protocol

The Nakamoto (ledger) consensus protocol can be abstracted in four main sub-routines. The main subroutine has two modes operation: regular and bootstrapping. In the bootstrapping mode, the node listens to the network for a certain period of time to collect sufficient number of blocks so that it determines a legitimate chain to build upon. The node will remain in this mode for a number

<sup>9</sup> A real-world mining pool or party of a certain hashing power can be thought of as a set of flat-model parties.

of rounds that is determined by a parameter  $\Delta_{\text{bootstr}}$ . When that time window passes the node will switch to regular mode.

During regular mode, the node’s state can be abstracted as a pair  $(st, \mathcal{C})$ . We follow previous work that abstracted the Bitcoin backbone [13] and we abstract considerations related to managing secret-keys and transaction issuance on behalf of the node. These would be incorporated in an input selection function  $I(\cdot)$  that acts and updates the node’s state  $st$ . We focus on the chain data structure  $\mathcal{C}$  that is defined as follows. Let  $G(\cdot)$  and  $H(\cdot)$  be cryptographic hash functions with output in  $\{0, 1\}^\kappa$ . A block with target  $T \in \mathbb{N}$  is a quadruple of the form  $B = \langle r, prev, x, ctr \rangle$  where  $prev \in \{0, 1\}^\kappa$ ,  $x \in \{0, 1\}^*$ , and  $r, ctr \in \mathbb{N}$  are such that they satisfy the predicate  $\text{validblock}^T(B)$  defined as  $(H(ctr, G(r, st, x)) < T) \wedge (ctr < 2^{32})$ . Observe that each block  $B$  is associated with a timestamp  $r$ .

A *blockchain*, or simply a *chain*  $\mathcal{C}$  is a (possibly empty) sequence of *blocks*; the rightmost block by convention is denoted by  $\text{head}(\mathcal{C})$  (note  $\text{head}(\varepsilon) = \varepsilon$ ) The blocks in a chain are connected in the sense that if  $B_i = \langle r_i, prev_i, x_i, ctr_i \rangle$  then  $B_{i+1} = \langle r_{i+1}, prev_{i+1}, x_{i+1}, ctr_{i+1} \rangle$  that satisfies  $prev_{i+1} = H(ctr_i, G(r_i, prev_i, x_i))$ . By convention  $B_0 = \varepsilon$  and  $prev_0 = \varepsilon$ . In practice,  $B_0$ , a.k.a. the “genesis” block, may be selected to be a valid block and furthermore contain some unpredictable string to ensure that no attacker could pre-mine blocks. Nevertheless, these features are not significant for our analysis and we ignore them.

Regarding the timestamp sequence of a chain, it should satisfy the condition:

$$r_{i+1} > \text{median}(r_i, \dots, r_{i-2k_{\text{med}}+2}), \text{ for } i > 2k_{\text{med}} - 1 \quad (1)$$

where  $k_{\text{med}}$  is a protocol parameter which in the case of Bitcoin’s parameterization it is set to 6 blocks.

At any point in the execution a node needs to determine the current time. We abstract this by a query `RequestTime` to the clock functionality, which responds with a reading that is within  $\Phi$  of the correct time. Note that, in practice, Bitcoin achieves that by querying the system time, the median time of its neighbors in the peer-to-peer network as well as the human operator if a substantial deviation exists between the first two readings.<sup>10</sup> Such considerations are abstracted in our modeling by the slack that is adversarially introduced in the `RequestTime` response from the clock functionality. In terms of determining the timestamp to use for the next block, the node should take into account the rule of the median of the past  $2k_{\text{med}} - 1$  blocks (see equation (1) above): the current time will be “pushed” forward to ensure that it is ahead of the median.<sup>11</sup>

We measure the length  $\text{len}(\mathcal{C})$  of a chain  $\mathcal{C}$  as the number of blocks it contains. We will use the notation  $\mathcal{C}^{\uparrow k}$  to denote the chain that results after “pruning”

<sup>10</sup> As stated in <https://github.com/bitcoin/bitcoin/blob/master/src/timedata.cpp>, “never go to sea with two chronometers; take one or three” (cf. *triple modular redundancy*).

<sup>11</sup> Refer to the source code, <https://github.com/bitcoin/bitcoin/blob/master/src/miner.cpp>, line 30: `nNewTime = std::max(pindexPrev->GetMedianTimePast()+1, GetAdjustedTime());`.

the  $k$  latest blocks. Moreover, if  $\mathcal{C}_1$  is a prefix of  $\mathcal{C}_2$  we write  $\mathcal{C}_1 \preceq \mathcal{C}_2$ . The data contained in the blocks of a chain, collected in a vector will be denoted by  $\mathbf{x}_{\mathcal{C}}$ . Similarly,  $\mathbf{r}_{\mathcal{C}}$  is the sequence of timestamps of  $\mathcal{C}$ .

Next, we describe how the appropriate target  $T$  is determined for each block. The initial target is a parameter of the protocol denoted by  $T_0$ . In the Bitcoin parameterization this is set to  $2^{224}$  and is the highest possible target. The target is updated every  $m$  blocks, which is another parameter of the protocol. In the Bitcoin parameterization it is set to 2016. The initial number of parties in the system, denoted by  $n_0$ , are assumed to be capable of producing  $m$  blocks in a length of time equal to  $\Delta_{\text{epoch}}$ . In Bitcoin’s parameterization  $\Delta_{\text{epoch}}$  corresponds to a length of time equal to 2 weeks. We will refer to this length of time as an “epoch.” The difficulty of each block is measured in terms of how much harder it is to produce a block compared to a block using the initial target  $T_0$ ; i.e., the blocks of the initial epoch have all difficulty 1. We take a slightly more general approach and define the difficulty of a block to be  $1/T$ . We will use the notation  $\text{diff}(\mathcal{C})$  to denote the sum of the difficulties of all blocks in chain  $\mathcal{C}$ .

The difficulty of the next block to be mined is determined by a function  $D(\cdot)$  that takes as input the sequence of timestamps corresponding a given chain  $\mathcal{C}$ . The function  $D(\cdot)$  parses the timestamps and identifies the last complete epoch of  $m$  blocks; if no such epoch is defined by  $\mathcal{C}$  the target is by definition equal  $T_0$ . Let  $r$  be the timestamp corresponding to last block of the last complete epoch,  $T$  its target and  $r_0$  the timestamp corresponding to the last block of the previous epoch (or the genesis block timestamp if such epoch does not exist). The function  $D(\cdot)$  returns as the next target the value  $T \cdot \frac{r-r_0}{\Delta_{\text{epoch}}}$ , unless  $\frac{r-r_0}{\Delta_{\text{epoch}}} \notin [\frac{1}{\tau}, \tau]$ , in which case it returns  $\tau \cdot T$ , when  $\frac{r-r_0}{\Delta_{\text{epoch}}} > \tau$ , or  $1/\tau \cdot T$ , otherwise. The parameter  $\tau$  is a “dampening factor” and in the case of Bitcoin’s parameterization it is set to 4.

**Blockchain properties.** The blockchain data structure’s two fundamental properties, adapted from [12, 13], are related to the Nakamoto consensus properties of Consistency and Liveness (see below).

The *common prefix* property, parameterized by a value  $k \in \mathbb{N}$ , considers an arbitrary environment and adversary, and holds as long as any two parties’ chains at two rounds have the earlier one subsumed in the former as long as  $k$  blocks are removed.

**Definition 2 (Common Prefix).** *The common prefix property  $Q_{\text{cp}}$  with parameter  $k \in \mathbb{N}$  states that for any two players  $P_1, P_2$  holding chains  $\mathcal{C}_1, \mathcal{C}_2$  at rounds  $r_1, r_2$ , with  $r_1 \leq r_2$ , it holds that  $\mathcal{C}_1^{[k]} \preceq \mathcal{C}_2$ .*

The second property, called *chain quality*, quantifies the honest-party contributions in terms of aggregate difficulty that are contained in a sufficiently long and continuous portion of a party’s chain. As a result, the property restricts the amount of difficulty (and hence number of blocks) contributed by the adversary to any sufficiently long segment of the chain.

**Definition 3 (Chain Quality).** *The chain quality property  $Q_{\text{cq}}$ , with parameters  $\mu \in \mathbb{R}$  and  $\ell \in \mathbb{N}$ , states that for any party  $P$  with chain  $\mathcal{C}$  in  $\text{VIEW}_{\Pi, A, Z}$ ,*

and any segment of that chain of difficulty  $d$  such that the first block of the segment was computed at least  $\ell$  rounds earlier than the last block, the blocks the honest parties have contributed in the segment have total difficulty at least  $\mu \cdot d$ .

**Nakamoto consensus properties.** As mentioned in Section 1, *Nakamoto consensus* (aka “ledger consensus”) is the problem where a set of servers (nodes, parties) operate continuously accepting inputs (“transactions”) and incorporate them in a public data structure called the *ledger*. More specifically, the problem is to maintain a ledger of transactions serialized in the form of a transaction sequence  $\mathcal{L}$ ; satisfying the following two properties [12, 13]. Below we make the distinction between  $\mathcal{L}$  and  $\tilde{\mathcal{L}}$ , with the first denoting the settled ledger in the view of the party, and the second denoting the settled ledger with a sequence of transactions appended that are still not settled in the view of the party. In the context of Nakamoto’s Bitcoin protocol, we note that  $\tilde{\mathcal{L}}$  will be the sequence of transactions defined by the chain  $\mathcal{C}$  held by the party, while  $\mathcal{L}$  will be the sequence of transactions defined by the prefix  $\mathcal{C}^{\lceil k}$ , where  $k$  is a security parameter.

- **Consistency** (cf. **Persistence** [12]): For any two honest parties  $P_1, P_2$ , reporting  $\mathcal{L}_1, \mathcal{L}_2$  at rounds  $r_1 \leq r_2$ , respectively, it holds that  $\mathcal{L}_1$  is a prefix of  $\mathcal{L}_2$ .
- **Liveness** (parameterized by  $u \in \mathbb{N}$ , the “wait time” parameter): If a transaction  $tx$  is provided to all honest parties for  $u$  consecutive rounds, then it holds that for any player  $P$ ,  $tx$  will be in  $\mathcal{L}$ .

We remark that the problem is a variant of the *state machine replication* problem [28] (see also [11]).

## 4 Nakamoto Consensus in Bounded-Delay Networks

In this section we present the full analysis and proofs of Nakamoto’s consensus protocol in the originally envisioned dynamic environment where parties—without synchronized clocks—come and go, resulting in the adjustment of blocks’ difficulty values. In particular, we formally specify (Section 4.2) the two conditions mentioned in Section 1 under which consistency and liveness of Nakamoto consensus can be shown.

### 4.1 Additional notation, definitions, and preliminary propositions

Following [13], our probability space is over all executions of length at most some polynomial in  $\kappa$  and  $\lambda$  and we denote by  $\mathbf{Pr}$  the probability measure of this space. Furthermore, let  $\mathcal{E}$  be a random variable taking values on this space and with a distribution induced by the random coins of all entities (adversary, environment, parties) and the random oracle.

If at round  $r$  exactly  $n$  parties query the oracle for target  $T$ , the probability at least one of them will succeed is

$$f(T, n) = 1 - (1 - pT)^n \leq pTn, \quad \text{where } p = 1/2^\kappa.$$

Note that  $\Delta_{\text{epoch}}$  and the initial target  $T_0$  implies in our model an initial estimate of the number of parties  $n_0$ ; specifically,  $n_0 = 2^\kappa m / (T_0 \Delta_{\text{epoch}})$ , i.e., the number of parties it takes to produce  $m$  blocks of difficulty  $1/T_0$  in time  $\Delta_{\text{epoch}}$ . We denote  $f_0 = f(T_0, n_0)$  and we drop the subscript from  $f_0$  and simply refer to it as  $f$ . Note the inequality

$$\frac{f(T, n)}{1 - f(T, n)} = \frac{1 - (1 - pT)^n}{(1 - pT)^n} = (1 - pT)^{-n} - 1 > (1 + pT)^n - 1 > pTn. \quad (2)$$

The first inequality is  $1/(1 - x) > 1 + x$  and the second is Bernoulli's.

We will next present some definitions which will allow us to introduce a few (“good”) properties. These properties are an intermediate step towards proving common prefix and chain quality, but are also interesting in their own. The next two definitions are about the notions of “good chain” and “good round.” The underlying notion of “goodness” is concerned with the targets that the honest parties are querying the random oracle for. At a round  $r$  of an execution the  $n_r$  honest parties might be querying the random oracle for various targets. We denote by  $T_r^{\min}$  and  $T_r^{\max}$  the minimum and maximum of those targets.

With respect to parameters that appear as “free” in the following definitions (such as  $\gamma, \Delta, \ell$ ), please refer to the next subsection.

**Definition 4.** – *Round  $r$  is good if  $f/2\gamma^2 \leq pn_r T_r^{\min}$  and  $pn_r T_r^{\max} \leq (1 + \delta)\gamma^2 f$ .*

- *Round  $r$  is a target-recalculation point of a chain  $\mathcal{C}$ , if  $\mathcal{C}$  has a block created in  $r$  and with height a multiple of  $m$ .*
- *A target-recalculation point  $r$  is good if the target  $T$  for the next block satisfies  $f/2\gamma \leq pn_r T \leq (1 + \delta)\gamma f$ .*
- *A chain is good if all its target-recalculation points are good.*
- *A chain is stale if for some  $u$  it does not contain an honest block computed in  $[u - \ell - 2\Delta, u]$ .*
- *A sequence of  $2K - 1$  consecutive blocks in a chain are called an honest median if there exists among them an honest block with timestamp the median of the timestamps of these blocks.*
- *The blocks between two consecutive target recalculation points  $u$  and  $v$  on a chain  $\mathcal{C}$  are an epoch of  $\mathcal{C}$  and the duration of the epoch is  $u - v$ .*

At a certain round of an execution, we would like to prove that the chain of every honest party has several desirable properties (along the notions just defined). This, however, entails a stronger statement in the following sense. At any given round there might exist chains which do not belong to any honest party (perhaps because the adversary kept them private), but have the potential to be adopted by one (i.e., have sufficient difficulty). With that in mind we define the following set of chains of a round  $r$ .

$$\mathcal{S}_r = \left\{ \mathcal{C} \in E_r \left| \begin{array}{l} (\mathcal{C} \text{ belongs to an honest party) or } (\exists \mathcal{C}' \in E_r \text{ that belongs} \\ \text{to an honest party and either } \text{diff}(\mathcal{C}) > \text{diff}(\mathcal{C}') \text{ or} \\ \text{diff}(\mathcal{C}) = \text{diff}(\mathcal{C}') \text{ and head}(\mathcal{C}) \text{ was computed} \\ \text{no later than head}(\mathcal{C}')) \end{array} \right. \right\},$$



where  $\mathcal{C} \in E_r$  means that  $\mathcal{C}$  exists and is valid at round  $r$ .

Next, we define a series of useful predicates with respect to such set of chains.

**Definition 5.** For a round  $r$ , let:

- $\text{GOODCHAINS}(r) \triangleq$  “For all  $u \leq r$ , every chain in  $\mathcal{S}_u$  is good.”
- $\text{GOODROUNDS}(r) \triangleq$  “All rounds  $u \leq r$  are good.”
- $\text{NOSTALECHAINS}(r) \triangleq$  “For all  $u \leq r$ , there are no stale chains in  $\mathcal{S}_u$ .”
- $\text{MEDIANTIME}(r) \triangleq$  “For all  $u \leq r$  and  $\mathcal{C} \in \mathcal{S}_u$ , there is an honest median computed in the last  $\lfloor \epsilon^2 m / f \rfloor$  rounds of any completed epoch in  $\mathcal{C}$ .”
- $\text{DURATION}(r) \triangleq$  “For all  $u \leq r$  and  $\mathcal{C} \in \mathcal{S}_u$ , the duration  $\Lambda$  of any epoch in  $\mathcal{C}$  satisfies  $\frac{1}{2(1+\delta)\gamma^2} \cdot \frac{m}{f} \leq \Lambda \leq 2(1+\delta)\gamma^2 \cdot \frac{m}{f}$ .”
- $\text{COMMONPREFIX}(r) \triangleq$  “For all  $u \leq r$  and  $\mathcal{C}, \mathcal{C}' \in \mathcal{S}_u$ ,  $\text{head}(\mathcal{C} \cap \mathcal{C}')$  was created after round  $u - \ell - 2\Delta$ .”

Our goal is to show that, with high probability, an execution satisfies the common prefix and chain quality properties. To fulfill this goal we will first focus on showing that the execution satisfies the predicates we defined above. In particular, we will argue first that none of these predicates can fail, assuming proper initialization. Towards that goal, we first specify a number of random variables.

**Random variables.** In our analysis, we will be interested in estimating the difficulty acquired by honest parties during a sequence of rounds. Their number at a round  $r$  is denoted  $n_r$  and define the real random variable  $D_r$  equal to the sum of the difficulties of all blocks computed by honest parties at round  $r$ . Also, define  $Y_r$  to equal the maximum difficulty among all blocks computed by honest parties at round  $r$ , and  $Q_r$  to equal  $Y_r$  when  $D_u = 0$  for all  $r < u < r + \Delta$  and 0 otherwise. We call a round  $r$  such that  $D_r > 0$  *successful* and one wherein  $Q_r > 0$  *isolated successful*. Regarding the adversary, let  $t_r$  denote the number of parties he controls at round  $r$  (equivalently, the number of random-oracle queries he can make at round  $r$ ). Note that  $n_r$  and  $t_r$  are determined by the environment at the beginning of round  $r$  and should conform to the  $(\gamma, s)$ -respecting definition (Definition 1). We wish to upper bound the difficulty he can acquire during a set  $J$  of queries. Looking ahead, to obtain a good upper bound that holds with high probability, we will need some upper bound on the difficulty of a single block. However, the adversary may query the oracle for arbitrarily low targets and may obtain blocks of arbitrarily high difficulty. The following definition will allow us to work around these technical obstacles.

Consider a set of consecutive adversarial queries  $J$  and note that the execution up to the first query in  $J$  determines the target associated with it. We denote this target by  $T(J)$  and say that  $T(J)$  is *associated* with  $J$ . We define  $A(J)$  and  $B(J)$  to be equal to the sum of the difficulties of all blocks computed by the adversary during queries in  $J$  for target at least  $T(J)/\tau$  and  $T(J)$ , respectively. That is, queries in  $J$  for targets less than  $T(J)/\tau$  (resp.  $T(J)$ ) do not contribute to  $A(J)$  (resp.  $B(J)$ ). While considering consecutive epochs of a

particular chain, the target can either increase by at most  $\tau$  (and  $B(J)$  will be appropriate), or decrease by at most  $\tau$  (and  $A(J)$  will be useful).

For a set of rounds  $S$  or queries  $J$  we write  $n(S) = \sum_{r \in S} n_r$  and similarly  $t(S)$ ,  $D(S)$ ,  $Q(S)$ . An interval of rounds (or queries) is a set of consecutive rounds and is denoted using bracket notation. For example, if  $u$  and  $v$  are two rounds such that  $u \leq v$  we write  $[u, v)$  for  $\{r : u \leq r < v\}$ .

Let  $\mathcal{E}_{r-1}$  fix the execution just before round  $r$ . In particular, a value  $E_{r-1}$  of  $\mathcal{E}_{r-1}$  determines the adversarial strategy and so determines the targets against which every party will query the oracle at round  $r$  and the number of parties  $n_r$  and  $t_r$ , but it does not determine  $D_r$  or  $Q_r$ . For an adversarial query  $j$  we will write  $\mathcal{E}_{j-1}$  for the execution just before this query.

**Proposition 1.** *Consider an execution  $E_{r-1}$  such that  $n_r = n$ ,  $T_r^{\max} = T^{\max}$ , and  $T_r^{\min} = T^{\min}$ . Then,*

$$\begin{aligned} [1 - f(T^{\max}, n)]pn &\leq \mathbf{E}[Y_r | \mathcal{E}_{r-1} = E_{r-1}] \leq \mathbf{E}[D_r | \mathcal{E}_{r-1} = E_{r-1}] = pn, \\ \mathbf{E}[Y_r^2 | \mathcal{E}_{r-1} = E_{r-1}] &\leq pn/T^{\min}, \quad \mathbf{var}[D_r | \mathcal{E}_{r-1} = E_{r-1}] \leq pn/T^{\min}. \end{aligned}$$

*Proof.* Suppose that the  $n$  honest parties at round  $r$  query for targets  $T_1, \dots, T_n$ . Observe that all these variables are determined by  $E_{r-1}$ . We have

$$\begin{aligned} \mathbf{E}[Y_r | \mathcal{E}_{r-1} = E_{r-1}] &= \sum_{i \in [n]} \frac{1}{T_i} \cdot \frac{T_i}{2^\kappa} \prod_{i < j} [1 - f(T_j, 1)] \geq \sum_{i \in [n]} p \prod_{j \in [n]} [1 - f(T_j, 1)] \\ &\geq \sum_{i \in [n]} p \prod_{j \in [n]} [1 - f(T^{\max}, 1)] = \sum_{i \in [n]} p[1 - f(T^{\max}, n)] = pn[1 - f(T^{\max}, n)], \end{aligned}$$

where the third inequality holds because  $f(T, n)$  is increasing in  $T$ . For the upper bound on variance,

$$\mathbf{var}[D_r | \mathcal{E}_{r-1} = E_{r-1}] \leq \sum_{i \in [n]} \frac{1}{T_i^2} \cdot \frac{T_i}{2^\kappa} = \sum_{i \in [n]} \frac{p}{T_i} \leq \frac{pn}{T^{\min}}$$

and  $\mathbf{E}[Y_r^2 | \mathcal{E}_{r-1} = E_{r-1}]$  is bounded alike.  $\square$

The following proposition collects a few useful inequalities that hold in a  $(\gamma, s)$ -respecting environment.

**Proposition 2.** *Let  $U$  be a set of at most  $s$  consecutive rounds in a  $(\gamma, s)$ -respecting environment and  $S \subseteq U$ .*

- (a) *For any  $n \in \{n_r : r \in U\}$ ,  $\frac{n}{\gamma} \leq \frac{n(S)}{|S|} \leq \gamma n$ .*
- (b)  *$n(U) \leq (1 + \frac{\gamma|U \setminus S|}{|S|})n(S)$ .*
- (c)  *$|S| \sum_{r \in S} (pn_r)^2 \leq \gamma (\sum_{r \in S} pn_r)^2$ .*

*Proof.* The first part is proved in [13] and is a direct consequence of the definition. For the second, note that  $n(U) = n(S) + n(U \setminus S)$ . By the first part we obtain that the greatest  $n \in \{n_r : r \in U \setminus S\}$  is at most  $\gamma n(S)/|S|$  and so  $n(U \setminus S) \leq |U \setminus S| \gamma n(S)/|S|$ . For the third, note that  $pn/\gamma \leq pn_r \leq \gamma pn$  for any  $r \in S$ . The inequality follows from Theorem 10, since  $A/G \leq (\gamma + 1/\gamma)/2 \leq \gamma$ .  $\square$

## 4.2 Parameters and Their Conditions

Our formalization of the protocol involves a number of parameters that have already appeared in the text above. With respect to the environment,

- $\delta$ : Advantage of honest parties<sup>12</sup> ( $t_r < (1 - \delta)n_r$  for all  $r$ );
- $(\gamma, s)$ : It restricts the fluctuation of the number of parties across rounds (Definition 1); we set  $s = \tau m/f$ .

An important parameter, which is a function of the protocol's initialization parameters  $n_0$  and  $T_0$ , is

- $f \in (0, 1)$ : The probability at least one honest party out of  $n_0$  computes a block for target  $T_0$ ; i.e.,  $f(T_0, n_0)$ .

The protocol strives to maintain the probability of a successful round as close  $f$  as possible. This is the task of the target recalculation function.

- $m \in \mathbb{N}$ : The length of an epoch in number of blocks;
- $\tau \geq 1$ : The dampening filter.

The value  $\tau m/f$  is the longest an epoch might take to complete and  $\gamma \geq 1$  is an estimate of how much the number of parties can change in such a time interval. The next two parameters are related to security.

- $\lambda$ : The security parameter;
- $\kappa = \Omega(\lambda)$ : The length of the range of the hash function.

To achieve security, we will argue concentration of several random variables. Furthermore, in any exponentially long (in the security parameters) execution bad events are bound to happen.

- $\epsilon$ : Quality of concentration of random variables;
- $L$ : The total number of rounds in the execution.

For  $L$  polynomially bounded in  $\lambda$ , our statements will fail with probability  $\text{poly}(\lambda)(2^{-\kappa} + e^{-\lambda})$ .

In our analysis we will study events over intervals of rounds. In order to achieve desired probability of error (in the order of  $e^{-\lambda}$ ) we will need to work with intervals of at least  $\ell$  rounds, where

$$\ell = \left\lceil \frac{4(1 + 3\epsilon)}{\epsilon^2 f [1 - (1 + \delta)\gamma^2 f]^{\Delta+1}} \cdot \max\{\Delta, \tau\} \cdot \gamma^3 \cdot \lambda \right\rceil, \quad \text{for } \epsilon \leq \frac{\delta}{16}. \quad (3)$$

For our analysis to go through, the above parameters should satisfy certain **conditions** which we now discuss. First, we will require a lower bound on the epoch length which incorporates  $\ell$ .

$$\frac{\epsilon^2 m}{f} \geq \max \left\{ \frac{\lambda^3 4^{k_{\text{med}}+6}}{\delta^4} + \ell + 2\Delta, 8(\Delta_{\text{fwd}} + \Phi), (\ell + 3\Delta) \cdot \frac{2^5(1 + \delta)^2 \gamma^6}{(1 - \epsilon)^3} \right\}. \quad (\text{C1})$$

<sup>12</sup> Note that we denote the number of honest parties at round  $r$  by  $n_r$  and the number of parties controlled by the adversary by  $t_r$ , so that the total number of parties is  $n_r + t_r$ . Although this is not standard, it simplifies several expressions and is also in agreement with notation in [13].

Second, the number of rounds  $\Delta$  that the adversary may delay messages relative to the block production rate should be upper bounded.

$$[1 - (1 + \delta)\gamma^2 f]^{2\Delta} \geq 1 - \epsilon. \quad (\text{C2})$$

Third, we will require in our analysis that the fluctuation rate of the parties is absorbed by the honest parties' advantage.

$$\gamma \leq 1 + \delta/8. \quad (\text{C3})$$

The time-out parameter for bootstrapping  $\Delta_{\text{bootstr}}$ , should allow enough time for the messages of a joining party to reach active parties and enough time for their response to come back. The drift between the clocks of two honest parties should not be greater than  $\Delta_{\text{fwd}}$ , otherwise one honest party might not accept a block computed by another. These amount to requiring  $\Delta_{\text{bootstr}} \geq 2\Delta$  and  $\Delta_{\text{fwd}} \geq 2\Phi$ .

### 4.3 Chain Growth Lemma

We now prove the Chain Growth Lemma. This lemma appears already in [12] in a model with fixed difficulty and fixed number of parties. In [17] the name ‘‘chain growth’’ appears for the first time, and where the authors explicitly state a Chain Growth Property. In [13], the lemma is proved in a synchronous model allowing variable difficulty and varying number of parties. Here we give a different proof that works in the dynamic bounded-delay model. The lemma provides a lower bound on the progress of the honest parties, which holds irrespective of any adversary.

**Lemma 1 (Chain Growth).** *Suppose that at round  $u$  of an execution  $E$  an honest party broadcasts a chain of difficulty  $d$ . Then, by round  $v$ , every honest party has received a chain of difficulty at least  $d + Q(S)$ , where  $S = [u + \Delta, v - \Delta]$ .*

*Proof.* If two blocks are obtained at rounds which are at distance at least  $\Delta$ , then we are certain that the later block increased the accumulated difficulty. To be precise, assume  $S^* \subseteq S$  is such that, for all  $i, j \in S^*$ ,  $|i - j| \geq \Delta$  and  $Y_i > 0$ . We argue that, by round  $v$ , every honest party has a chain of difficulty at least

$$d + \sum_{r \in S^*} Y_r \geq d + \sum_{r \in S} Q_r.$$

Observe first that every honest party will receive the chain of difficulty  $d$  by round  $u + \Delta$  and so the first block obtained in  $S^*$  extends a chain of weight at least  $d$ . Next, note that if a block obtained in  $S^*$  is the head of a chain of weight at least  $d'$ , then the next block in  $S^*$  extends a chain of weight at least  $d'$ .  $\square$

### 4.4 Typical Executions

We now define formally our notion of *typical* executions. Intuitively, the idea that this definition captures is as follows. Suppose that we examine a certain

execution  $E$ . Note that at each round of  $E$  the parties perform Bernoulli trials with success probabilities possibly affected by the adversary. Given the execution, these trials are determined and we may calculate the expected progress the parties make given the corresponding probabilities. We then compare this value to the actual progress and if the difference is reasonable we declare  $E$  *typical*. Note, however, that considering this difference by itself will not always suffice, because the variance of the process might be too high. Our definition, in view of Theorem 8, says that either the variance is high with respect to the set of rounds we are considering, or the parties have made progress during these rounds as expected.

Beyond the behavior of random variables described above, a typical execution will also be characterized by the absence of a number of bad events about the underlying hash function  $H(\cdot)$  which is used in proofs of work and is modeled as a random oracle. The bad events that are of concern to us are defined as follows (recall that a block's creation time is the round that it has been successfully produced by a query to the random oracle either by the adversary or an honest party).

**Definition 6.** *An insertion occurs when, given a chain  $\mathcal{C}$  with two consecutive blocks  $B$  and  $B'$ , a block  $B^*$  created after  $B'$  is such that  $B, B^*, B'$  form three consecutive blocks of a valid chain. A copy occurs if the same block exists in two different positions. A prediction occurs when a block extends one with later creation time.*

Given the above we are now ready to specify what is a typical execution.

**Definition 7 (Typical execution).** *An execution  $E$  is typical if the following hold.*

(a) *For any set  $S$  of at least  $\ell$  consecutive good rounds,*

$$(1 - \epsilon)[1 - (1 + \delta)\gamma^2 f]^\Delta pn(S) < Q(S) \leq D(S) < (1 + \epsilon)pn(S).$$

(b) *For any set  $J$  of consecutive adversarial queries and  $\alpha(J) = 2(\frac{1}{\epsilon} + \frac{1}{3})\lambda/T(J)$ ,*

$$A(J) < p|J| + \max\{cp|J|, \tau\alpha(J)\} \quad \text{and} \quad B(J) < p|J| + \max\{cp|J|, \alpha(J)\}.$$

(c) *No insertions, no copies, and no predictions occurred in  $E$ .*

We will be interested in comparing the computational power of the adversary against that of the honest parties in a set of consecutive rounds  $S$ . However, in the bounded-delay model with delay  $\Delta$ , the adversary can mute the honest parties for the final  $\Delta$  rounds. Calculations summarized in the following lemma will be used repeatedly.

**Lemma 2.** *Consider a typical execution in a  $(\gamma, s)$ -respecting environment. Let  $S = \{r : u \leq r \leq v\}$  be a set of at least  $\ell$  consecutive good rounds and  $J$  the set of adversarial queries in  $U = \{r : u - \Delta \leq r \leq v + \Delta\}$ .*

- (a)  $(1 + \epsilon)p|J| \leq Q(S) \leq D(U) < (1 + 4\epsilon)Q(S)$ .
- (b)  $T(J)A(J) < \frac{(1-\epsilon)^3}{32(1+\delta)^2\gamma^9} \cdot \epsilon^2 m$  or  $A(J) < (1 + \epsilon)p|J|$ ; similarly  
 $\tau T(J)B(J) < \frac{(1-\epsilon)^3}{32(1+\delta)^2\gamma^9} \cdot \epsilon^2 m$  or  $B(J) < (1 + \epsilon)p|J|$ .
- (c) If  $w$  is a good round such that  $|w - r| \leq s$  for any  $r \in S$ , then  $Q(S) > (1 - \epsilon)[1 - (1 + \delta)\gamma^2 f]^\Delta |S|pn_w/\gamma$ . If in addition  $pn_w T(J) \geq f/2\gamma^2$ , then  $A(J) < (1 - \delta + 3\epsilon)Q(S)$ .
- (d) In at most  $\ell + 2\Delta$  rounds can be produced at most  $\frac{(1-\epsilon)^2}{8(1+\delta)\gamma^4} \cdot \epsilon^2 m$  blocks of a good chain.

The main result of this section is that almost all polynomially bounded in  $\kappa$  and  $\lambda$  executions are typical.

**Theorem 1.** *Assuming the ITM system  $(Z, C)$  runs for  $L$  steps, the probability of the event “ $\mathcal{E}$  is not typical” is bounded by  $O(L^2)(e^{-\lambda} + 2^{-\kappa})$ .*

#### 4.5 Hot-Hand Executions

In this section we define a new class of executions: *hot-hand executions*. The structure such executions are required to have will allow us to show that an honest median (cf. Definition 4) exists sufficiently often. At a high level, a hot-hand execution contains special streaks of honest successes. What is special about these streaks is the structure of the surrounding rounds, which is such that the honest blocks that correspond to the streak are guaranteed to form an honest median.

We define a sequence  $\{(V_i, R_i) : i > 0\}$  of random variables taking values in  $\mathbb{Z} \times \mathbb{N}$ , with respect to a sequence of rounds  $r_1, r_2, \dots$  and a target  $T$ . A round  $r$  is called *strongly-isolated successful* if an honest query was successful for target  $T$  in round  $r$  and no other query was successful for target  $T$  in the interval of rounds  $(r - \Delta, r + \Delta)$ . Set  $R_1 = r_1$  and for  $i > 0$  set  $R_{i+1} = r_{j+1}$  where  $r_j \geq R_i$  is the least round such that either an adversarial block was created in  $r_j$  or  $r_j - \Delta + 1$  was isolated successful. In the former case  $V_i$  is minus the number of blocks acquired by the adversary during round  $r_j$  and in the latter it is equal to one. For an interval of rounds  $S = [u, v]$ , we denote  $\sum V_i$ , where the sum is over  $\{i : u \leq R_i \leq v\}$ , as  $V(S)$  when the sequence  $\{(V_i, R_i) : i > 0\}$  is defined with respect to the sequence of rounds  $u, u + 1, \dots, v$ .

**Definition 8.** *An interval that contains at least  $k_{\text{med}}$  strongly-isolated successful rounds and no other honest or adversarial blocks were computed in it is called a streak. A streak  $[u, v]$  such that  $V(S) \geq 0$  for any  $S = (v, v']$  and  $V(S) \geq 0$  for any  $S = [u', u)$  is a winning streak.*

For a given execution  $E$  we associate with each set  $S = [u, v]$  of at most  $s$  consecutive rounds the target  $T$  of the first honest query in  $S$ . We want to say that  $T$  is good if  $E_u$  forces  $T$  to be good for each round in  $S$ . Formally, let  $n'$  be the maximum number of honest parties in  $U = (v - s, u]$ :  $n' = \max_{r \in U} n_r$ . Note that  $E_u$  determines  $n'$  and that, since  $|U \cup S| = s$ , a  $(\gamma, s)$ -respecting

environment forces  $n'/\gamma \leq n_r \leq \gamma n'$  for each  $r \in S$ . We say  $T$  is *good* for  $S$  in  $E$ , if  $pn'T \leq (1+\delta)\gamma f$ . Note that this implies  $pn_r T \leq (1+\delta)\gamma^2 f$  for each  $r \in S$ .

**Definition 9 (Hot-hand execution).** *An execution  $E$  is hot-hand if for any interval  $S$  of at least  $\delta^{-4}4^{k_{\text{med}}+6}\lambda^3$  rounds such that the associated target  $T$  is good for  $S$  in  $E$ , there is a winning streak in  $S$  with respect to  $T$ .*

We wish to show that an execution is hot-hand with overwhelming probability. Since we are studying a small set of rounds each time, we are going to absorb the fluctuation  $\gamma$  in the number of parties in the advantage  $\delta$  of the honest parties. Specifically, recall the definition of  $n'$  in the paragraph before Definition 9 and set  $n = \lceil n'/\gamma \rceil$ . In a  $(\gamma, s)$ -respecting environment, the number of honest parties in  $S$  may fluctuate between  $n$  and  $\gamma^2 n$  (because  $\gamma^2 n \geq \gamma n' \geq \lfloor \gamma n' \rfloor$ ). Our plan is to reduce the question of the existence of a winning streak in  $S$  to a question in a static execution with  $n$  honest parties and a non-adaptive adversary controlling a fixed number of  $t < (1 - \delta/2)n$  parties. This is going to work because, for  $\gamma \leq 1 + \delta/8$ ,

$$(\gamma^2 - 1)n + (1 - \delta)\gamma^2 n \leq [(2 - \delta)(1 + \delta/8)^2 - 1]n < (1 - \delta/2)n.$$

Thus, we may handle the fluctuation of parties above  $n$ , by allowing the adversary to control all of the excess parties.

We make this argument formal via a coupling argument. We show that the probability that an interval of rounds is a winning streak in this static setting is at most that in the dynamic setting where the honest parties fluctuate between  $n$  and  $\gamma^2 n$  against an adaptive adversary. To prove this, we generate both distributions by starting with the maximum number  $\lfloor (2 - \delta/2)\gamma^2 n \rfloor$  of oracle outputs per round. In the dynamic setting, the strategy of the adaptive adversary determines in the beginning of each round how many of these outputs will be discarded, while in the static setting this number is fixed. Note that in the dynamic setting there are at least as many honest parties as in the static one. Furthermore, we distribute this excessive number of at most  $(\gamma^2 - 1)n$  honest queries to the static adversary. Under this construction, every honest query in the static setting is also an honest query in the dynamic. Conversely, every query in the dynamic setting is a query in the static one. The following properties follow.

- Every successful round in the static setting is also successful in the dynamic.
- If there are no successful queries in a round of the static setting, then the same holds for this round in the dynamic setting.

We observe now that if an interval  $[u, v]$  is a winning streak for a given sequence of oracle outputs in the static setting, then this interval is also a winning streak for this sequence in the dynamic setting. Indeed, it follows directly from the above properties that any isolated successful round  $r$  in the static setting is still isolated in the dynamic one. Now consider an interval  $S = [u', u)$ . If  $V(S)$  is non-negative in the static setting, then it is also non-negative in the dynamic setting, as only more honest queries are introduced.

**Winning Streaks in the Static Setting.** We now proceed with the analysis in the static setting outlined above.

**Lemma 3.** *For any round  $u$ , the probability that  $V(S) \geq 0$  for all  $S = [u, v]$  is at least  $\delta/8$ .*

*Proof.* This follows by Theorem 12 applied to the sequence  $(V_i : i > 0)$ . We need to show that  $\mathbf{E}[V_i] \geq \delta/8$  for any  $i > 0$ .

For each nonnegative integer  $k$  set  $q_k = \mathbf{Pr}[V_i = -k]$  and  $q = \mathbf{Pr}[V_i = 1]$ . Setting  $p_T = pT$  and recalling  $\sum_{k \geq 0} x^k = 1/(1-x)$ ,

$$\begin{aligned} q &= \sum_{r \geq 0} (1-p_T)^{(n+t)r} \cdot \frac{p_T^n}{1-p_T} \cdot (1-p_T)^{(n+t)(2\Delta-1)} \\ &= \frac{1}{1-(1-p_T)^{(n+t)}} \cdot \frac{p_T^n}{1-p_T} \cdot (1-p_T)^{(n+t)(2\Delta-1)}. \end{aligned} \quad (4)$$

Under the hypothesis that  $T$  is good it holds  $p_T n \leq (1+\delta)\gamma f$ . Using this, Bernoulli's inequality, and Condition C2 we obtain the lower bound

$$q > \frac{p_T n (1-p_T)^{2(n+t)\Delta}}{p_T (n+t)} > \frac{[1-(1+\delta)\gamma f]^{(4-\delta)\Delta}}{2-\delta/2} > \frac{(1-\epsilon)^2}{2-\delta/2} > \frac{2-4\epsilon}{4-\delta}.$$

Since  $\mathbf{E}[V_i] = q - \sum_{k>0} k q_k$ , we now turn to the sum.

$$\begin{aligned} \sum_{k>0} k q_k &= \sum_{k>0} k \sum_{r \geq 0} (1-p_T)^{(n+t)r} \binom{t}{k} p_T^k (1-p_T)^{n+t-k} \\ &= (1-p_T)^n \sum_{r \geq 0} (1-p_T)^{(n+t)r} \sum_{k>0} k \binom{t}{k} p_T^k (1-p_T)^{t-k} \\ &= p_T t (1-p_T)^n \sum_{r \geq 0} (1-p_T)^{(n+t)r} = \frac{p_T t (1-p_T)^n}{1-(1-p_T)^{(n+t)}}. \end{aligned}$$

The third equality holds because the last sum is the expected value of a binomial and equals  $p_T t$  and the last equality by the identity  $\sum_{k \geq 0} x^k = 1/(1-x)$ . Next, by inequality (2),  $p_T(n+t)(1-p_T)^{n+t} < 1 - (1-p_T)^{n+t}$ ; thus,

$$\begin{aligned} \sum_{k>0} k q_k &< \frac{p_T t (1-p_T)^n}{p_T (n+t) (1-p_T)^{(n+t)}} < \frac{2-\delta}{4-\delta} \cdot \frac{1}{(1-p_T)^n} < \frac{2-\delta}{4-\delta} \cdot \frac{1}{1-p_T n} \\ &< \frac{2-\delta}{4-\delta} \cdot \frac{1}{1-(1+\delta)\gamma f} < \frac{2-\delta}{4-\delta} \cdot \frac{1}{\sqrt{1-\epsilon}} < \frac{(2-\delta)(1+\epsilon)}{4-\delta}. \end{aligned}$$

For the second inequality we use  $t < (1-\delta/2)n$ , the third is Bernoulli's, the fourth is  $p_T n \leq (1+\delta)\gamma^2 f$ , the fifth Condition C2. Finally, since  $\epsilon \leq \delta/16$ ,

$$\mathbf{E}[V_i] > \frac{2-\delta/4}{4-\delta} - \frac{2-\delta+\delta/8}{4-\delta} > \frac{\delta}{8}.$$

□



Let  $X_r$  and  $Z_r$  denote the number of successful queries by honest parties and the adversary in round  $r$  respectively. Consider an interval  $S$  and for  $(X, Z) = (X_r, Z_r : r \in S)$  let  $h(X, Z)$  be equal to the number of winning streaks beginning at a round in  $S$ . We claim that

$$\mathbf{E}[h] > \delta^2 |S| / 2^{k_{\text{med}}+6}. \quad (5)$$

To verify this, observe that since  $q > 1/2$  the probability a streak begins at a round  $r$  is at least  $1/2^{k_{\text{med}}}$ . Irrespectively of the round  $r$ , by Lemma 3, the probability that  $V([u, r]) \geq 0$  for all  $u$  is at least  $\delta/8$ . Similarly, for any round  $v$  that the final block of the streak is computed, the probability that  $V([v, u]) \geq 0$  is at least  $\delta/4$ . Since these events are independent, the bound follows.

We will now show that  $h(X, Z)$  is concentrated around its expected value. To that end, it will be convenient to assume that each  $X_r$  and  $Z_r$  never surpass  $\lambda$ . The probability more than  $\lambda$  queries out of  $n$  are successful can be bounded as in [22, Theorem 4.4] by  $(ep_T n / \lambda)^\lambda < e^{-\lambda}$  (note that Condition C2 implies  $(1 + \delta)\gamma^2 f \leq \epsilon \leq \delta/16$  and so  $\lambda \geq 1 \geq 16p_T n > e^2 p_T n$ ).

**Lemma 4.** *Assuming that neither the honest parties nor the adversary acquired more than  $k$  blocks in a single round of an interval  $S$ , then  $f$  is  $K$ -Lipschitz over  $S$  for  $K = \max\{k_{\text{med}} + 2, 2\lceil k/k_{\text{med}} \rceil\}$*

*Proof.* We are interested in the maximum value of  $|h(x, z) - h(x', z')|$  over  $(x, z), (x', z') \in \{0, 1\}^{|S|}$  that differ only on the  $i$ -th position. Consider winning streaks  $W = [u, v]$  and  $W' = [u', v']$  and a round  $r$ . We say  $W'$  lies between  $W$  and  $r$  if  $v < u'$  and  $v' < r$ . Suppose  $x_i > x'_i$  and  $x_i > 1$ . We claim that any winning streak  $W$  of  $(x, z)$  with the property that at least  $k/k_{\text{med}}$  winning streaks lie between  $W$  and  $i$  is also a winning streak of  $(x', z')$ . This is because we can assume  $|x_i - x'_i| \leq k$  and so the winning streaks in between counter-balance the difference  $|x_i - x'_i|$ . Note also that  $x_i > 1$  implies that the blocks computed in round  $i$  do not themselves belong to a streak. It follows that even if  $x_i = k$  and  $x'_i = 0$ , the winning streaks that are “lost” in  $(x', z')$  are at most  $\lceil k/k_{\text{med}} \rceil$  to the right of  $i$  and at most the same number on the left. In the case  $x_i = 1$  and  $x'_i = 0$ , these streaks can be at most 2. However, the block computed at  $i$  can belong to at most  $k_{\text{med}}$  streaks. The statement follows.  $\square$

**Theorem 2.** *The probability an execution is not hot hand is  $\text{poly}(\lambda) \cdot e^{-\lambda}$ .*

*Proof.* We apply Theorem 9 to the sequences of random variables  $(X_j, Z_j)_{j \in [N]}$  that correspond to  $N = \lceil \delta^{-4} 4^{k_{\text{med}}+6} \lambda^3 \rceil$  rounds of a static execution with target  $T$ ,  $n$  honest parties,  $t < (1 - \delta/2)n$  adversarial parties, and  $\Gamma_j$  the event that neither the honest parties nor the adversary acquired more than  $\lambda$  blocks in round  $j$ . We apply the theorem with  $c_j = \lambda$ ,  $d = N$ ,  $\Pr[(X, Z) \notin \Gamma] \leq N e^{-\lambda}$  (as discussed before stating Lemma 4), and  $t = \delta^2 N / 2^{k_{\text{med}}+6} - N e^{-\lambda}$ . We obtain

$$\Pr[h(x) \leq 0] \leq \exp\left\{-\frac{\delta^4 N - O(N^2)e^{-\lambda}}{\lambda^2 4^{k_{\text{med}}+6}}\right\} + N^2 e^{-\lambda} = \text{poly}(\lambda) \cdot e^{-\lambda},$$

where the last equality uses Condition C1. We finish the proof with a union bound over the relevant intervals.  $\square$

**Towards the existence of an honest median.** We now prove a couple of lemmas that will be useful in the forthcoming analysis in proving the existence of an honest median. These lemmas describe what happens when an honest block is orphaned and will be useful in arguing that blocks in winning streaks which are deep in an epoch cannot be orphaned.

**Lemma 5.** *Suppose that for round  $v$  there exist chains  $\mathcal{C}$  and  $\mathcal{C}'$  in  $\mathcal{S}_v$  such that  $\mathcal{C} \setminus (\mathcal{C} \cap \mathcal{C}')$  contains an honest block computed in round  $r$ . Let  $u$  be the round that the last honest block on  $\mathcal{C} \cap \mathcal{C}'$  was computed. Then, for the interval  $S = (u, v]$  containing  $r$ ,  $V(S) \leq 0$ .*

*Proof.* Let  $\mathcal{B}$  be the set of honest blocks that contributed to  $V(S)$ . We prove  $V(S) \leq 0$  by exhibiting the existence of a set of adversarial blocks  $\mathcal{B}'$  computed in  $S$  such that  $\{d \in B : B \in \mathcal{B}\} \subseteq \{d \in B : B \in \mathcal{B}'\}$ .

Consider a block  $B \in \mathcal{B}$  extending a chain  $\mathcal{C}^*$  and let  $d = \text{diff}(\mathcal{C}^*B)$ . If  $d \leq \text{diff}(\mathcal{C} \cap \mathcal{C}')$ , let  $B'$  be the block of  $\mathcal{C} \cap \mathcal{C}'$  containing  $d$ . Such a block clearly exists on  $\mathcal{C} \cap \mathcal{C}'$  and was computed after round  $u$  by the adversary (due the definition of  $u$ ). If  $d > \text{diff}(\mathcal{C} \cap \mathcal{C}')$ , note that by the definition of an isolated successful round there is a unique  $B \in \mathcal{B}$  such that  $d \in B$ . Since  $\mathcal{C}$  and  $\mathcal{C}'$  are in  $\mathcal{S}_v$ , they are at least as heavy as  $\mathcal{C}^*$ . It follows that there is  $B' \notin \mathcal{B}$  either on  $\mathcal{C}$  or on  $\mathcal{C}'$  that contains  $d$ .  $\square$

**Lemma 6.** *Suppose a block was computed by an honest party at round  $r$  and does not belong to the chain of an honest party at round  $r' \geq r + \Delta$ . Then  $r$  is contained in an interval  $S$  such that  $V(S) \leq 0$ .*

*Proof.* Let  $B$  be the block computed at round  $r$  by an honest party  $P$ . We are going to apply the previous lemma for a round  $v$  and chains  $\mathcal{C}, \mathcal{C}' \in \mathcal{S}_v$  that we will define according to the following cases. If  $P$  adopted a chain  $\mathcal{C}'$  not containing  $B$ , then let  $v$  be the previous round. Otherwise, let  $v$  be the least round not earlier than  $r + \Delta$  such that an honest party  $P'$  does not have  $B$  in its chain  $\mathcal{C}'$ . Let  $\mathcal{C}$  be the chain of  $P$  at round  $v$ . In both cases,  $\mathcal{C}$  and  $\mathcal{C}'$  belong in  $\mathcal{S}_v$  and the previous lemma supplies the set  $S$ .  $\square$

## 4.6 Properties of Typical and Hot-Hand Executions

In this subsection we study in detail the validity of the predicates of Definition 5 over the space of typical and hot-hand executions in a  $(\gamma, s)$ -respecting environment. All statements in this subsection assume a  $(\gamma, s)$ -respecting environment for  $s \geq 2(1 + \delta)\gamma^2 m/f$ . Furthermore, the Conditions are assumed to hold for the initialization parameters  $n_0$  and  $T_0$ .

Our first lemma says that the adversary cannot maintain a chain by himself for too long. The reason is that the honest parties will progress faster and his blocks will be orphaned.

**Lemma 7.** *In a typical execution and a  $(\gamma, s)$ -respecting environment*

$$\text{GOODROUNDS}(r - 1) \implies \text{NOSTALECHAINS}(r).$$

*Proof.* Suppose—towards a contradiction— $\mathcal{C} \in \mathcal{S}_r$  and has not been extended by an honest party for at least  $\ell + 2\Delta$  rounds and  $r$  is the least round with this property. Let  $B$  be the last block of  $\mathcal{C}$  computed by honest parties at a round  $w$  (possibly  $w = 0$  and  $B$  the genesis). Set  $S = [w + \Delta, r - \Delta]$  and  $U = [w, r]$ . Note that by our assumption  $|S| \geq \ell$ . Suppose that the blocks of  $\mathcal{C}$  after  $B$  span  $k$  epochs with corresponding targets  $T_1, \dots, T_k$ . For  $i \in [k]$  let  $m_i$  be the number of blocks with target  $T_i$  and set  $M = m_1 + \dots + m_k$  and  $d = m_1/T_1 + \dots + m_k/T_k$ . Our plan is to contradict the assumption that  $\mathcal{C} \in \mathcal{S}_r$  by showing that all chains in  $\mathcal{S}_r$  have more difficulty than  $\mathcal{C}$ . By Chain-Growth (Lemma 1), all the honest parties have advanced (in difficulty) during the rounds in  $U$  by  $Q(S)$ . Therefore, to reach a contradiction it suffices to show that  $d < Q(S)$ .

When  $k > 2$  we may partition these  $M$  blocks into  $k - 1$  parts so that each part has the following properties: (1) it contains at most one target-recalculation point, and (2) it contains at least  $m/2$  blocks. For each  $i \in [1, k]$ , let  $j_i \in J$  be the index of the query during which the first block of the  $i$ -th part was computed and set  $J_i = [j_i, j_{i+1})$  (Definition 7(c) assures  $j_i < j_{i+1}$ ). We claim

$$d = \sum_{i=1}^k \frac{m_i}{T_i} < \sum_{i=1}^{k-1} (1 + \epsilon)p|J_i| \leq (1 + \epsilon)p|J| \leq Q(S).$$

For the first inequality, consider part  $i$ . We have  $T_i = T(J_i)$  and—because of the first property of the partition—two possible cases for  $T_{i+1}$ : either  $T_i \leq T_{i+1} \leq \tau T_i$  or  $T_i/\tau \leq T_{i+1} \leq T_i$ . In the first case, the difficulty of the blocks acquired in  $J_i$  is at most  $B(J_i)$  and their number at most  $\tau T_i B(J_i)$ . In the second case, the difficulty of the blocks acquired in  $J_i$  is at most  $A(J_i)$  and their number at most  $T_i A(J_i)$ . In either case, since the adversary acquired at least  $m/2$  blocks in  $J_i$ , the desired bound follows from Lemma 2(b). The final inequality is Lemma 2(a).

If  $k \leq 2$ , let  $J$  denote the queries in  $U$  starting from the first adversarial query attempting to extend  $B$ . Then,  $T_1 = T(J)$  and  $T_2 \geq T(J)/\tau$ ; thus,  $d \leq A(J)$ . If  $A(J) < (1 + \epsilon)p|J|$ , then  $d < Q(S)$  is obtained by Lemma 2(a). Otherwise,

$$A(J) < p|J| + \tau\alpha(J) < \left(\frac{1}{\epsilon} + 1\right)\tau\alpha(J) = 2\left(\frac{1}{\epsilon} + 1\right)\left(\frac{1}{\epsilon} + \frac{\epsilon}{3}\right)\tau\lambda/T(J),$$

where we used Definition 7 and the assumption  $A(J) \geq (1 + \epsilon)p|J|$ . Consider only the first  $\ell$  rounds in  $S$ . In a  $(\gamma, s)$ -respecting environment,  $pn(S) \geq pn_w \ell/\gamma$ . Furthermore, since  $w < r$ ,  $w$  is a good round and so  $pn_w T_1 \geq f/(2\gamma^2)$ . Putting these together,  $pn(S) \geq \ell f/(2\gamma^3 T_1)$ . By Definition 7 and the value of  $\ell$ ,

$$Q(S) > \frac{(1 - \epsilon)[1 - (1 + \delta)\gamma^2 f]^\Delta f \ell}{2\gamma^3 T(J)} \geq \frac{2(1 - \epsilon)(1 + 3\epsilon)\tau\lambda}{\epsilon^2 T(J)}.$$

Using the inequality for  $A(J)$  above and  $\epsilon \leq 1/8$ , we arrive at our desired contradiction  $d < Q(S)$ .  $\square$

The following lemma says that two “longest” chains cannot diverge for too long. We say below that  $d \in \mathbb{R}$  is *contained* in a block  $B$  and write  $d \in B$ , when  $B$  extends a chain  $\mathcal{C}$  and  $\text{diff}(\mathcal{C}) < d \leq \text{diff}(\mathcal{C}B)$ .

**Lemma 8.** *In a typical execution and a  $(\gamma, s)$ -respecting environment*

$$\text{GOODROUNDS}(r-1) \implies \text{COMMONPREFIX}(r).$$

*Proof.* Suppose  $\text{head}(\mathcal{C} \cap \mathcal{C}')$  was created in round  $v$  and let  $u \leq v$  be the greatest round in which an honest party computed a block on  $\mathcal{C} \cap \mathcal{C}'$ . Let  $U = [u, r]$ ,  $S = [u + \Delta, r - \Delta]$ , and let  $J$  denote the adversarial queries that correspond to the rounds in  $U$ . We claim that, if  $r - v \geq \ell + 2\Delta$ , then

$$2Q(S) \leq D(U) + A(J).$$

Let us first verify that this contradicts Lemma 2. First, if  $|S| \geq \ell$ , then by Lemma 2(a) it holds  $D(U) < (1 + 4\epsilon)Q(S)$ . Next, Lemma 7 implies that neither  $\mathcal{C}$  nor  $\mathcal{C}'$  is stale. This allows us to apply Lemma 2(c) and obtain  $A(J) < (1 - \delta + 4\epsilon)Q(S)$ . Putting these together with Condition C2 we obtain  $D(U) + A(J) < (2 - \delta + 8\epsilon)Q(S) < 2Q(S)$ .

Towards proving the claim, associate with each  $w \in S$  such that  $Q_w > 0$  an arbitrary honest block that is computed at round  $w$  for difficulty  $Q_w$ . Let  $\mathcal{B}$  be the set of these blocks and note that their difficulties sum to  $Q(S)$ . We argue the existence of a set of blocks  $\mathcal{B}'$  computed in  $U$  such that  $\mathcal{B} \cap \mathcal{B}' = \emptyset$  and  $\{d \in \mathcal{B} : B \in \mathcal{B}\} \subseteq \{d \in \mathcal{B} : B \in \mathcal{B}'\}$ . This suffices, because each block in  $\mathcal{B}'$  contributes either to  $D(U) - Q(S)$  or to  $A(J)$  and so  $Q(S) \leq D(U) - Q(S) + A(J)$ .

Consider a block  $B \in \mathcal{B}$  extending a chain  $\mathcal{C}^*$  and let  $d = \text{diff}(\mathcal{C}^*B)$ . If  $d \leq \text{diff}(\mathcal{C} \cap \mathcal{C}')$  (note that  $u < v$  in this case and  $\text{head}(\mathcal{C} \cap \mathcal{C}')$  is adversarial), let  $B'$  be the block of  $\mathcal{C} \cap \mathcal{C}'$  containing  $d$ . Such a block clearly exists and was computed after round  $u$ . Furthermore,  $B' \notin \mathcal{B}$ , since  $B'$  was computed by the adversary. If  $d > \text{diff}(\mathcal{C} \cap \mathcal{C}')$ , note that there is a unique  $B \in \mathcal{B}$  such that  $d \in B$  (recall the argument in Chain Growth Lemma 1). Since  $B$  cannot simultaneously be on  $\mathcal{C}$  and  $\mathcal{C}'$ , there is a  $B' \notin \mathcal{B}$  either on  $\mathcal{C}$  or on  $\mathcal{C}'$  that contains  $d$ .  $\square$

**Lemma 9.** *In a typical execution and a  $(\gamma, s)$ -respecting environment*

$$\text{GOODROUNDS}(r-1) \wedge \text{GOODCHAINS}(r-1) \implies \text{DURATION}(r).$$

*Proof.* Assume—towards a contradiction—that  $\text{DURATION}(r)$  is false. Then, there exists a  $w \leq r$  and a chain  $\mathcal{C} \in \mathcal{S}_w$  with an epoch of target  $T$  and duration  $\Lambda$  that does not satisfy

$$\frac{1}{2(1+\delta)\gamma^2} \cdot \frac{m}{f} \leq \Lambda \leq 2(1+\delta)\gamma^2 \cdot \frac{m}{f}.$$

We consider the earliest epoch for which one of these bounds on  $\Lambda$  fails.

For the upper bound, Lemma 7 implies the existence of two honest blocks in this epoch computed at least  $\Lambda - 2\ell - 4\Delta$  rounds apart. Let  $u$  and  $v$  be these rounds and define  $S = [u, v]$ . Assuming  $\Lambda > 2(1+\delta)\gamma^2 m/f$ , Condition C1 implies  $|S| \geq 2(1+\delta)(1-\epsilon)\gamma^2 m/f$ . Using this bound and our hypothesis that the rounds in  $S$  are good in Definition 7,

$$Q(S) > (1-\epsilon)[1 - (1+\delta)\gamma^2 f]^\Delta \cdot \frac{f|S|}{2\gamma^2 T} \geq (1+\delta)(1-\epsilon)^3 \cdot \frac{m}{T} > \frac{m}{T},$$

where we used Condition C2 for the last inequality. This contradicts Chain Growth, since the honest parties at round  $v$  already have more than  $m/T$  difficulty on top of  $u$ .

To prove the lower bound we are going to argue that even if the honest parties and the adversary join forces they still cannot obtain  $m$  blocks. Let  $u$  and  $v$  be the target-recalculation points of the epoch. Define  $S = [u, v]$  and  $J$  the set of queries available to the adversary during the rounds in  $S$  starting with the first query for target  $T$  (so that  $T(J) = T$ ). Without loss of generality, assume  $S$  has size exactly  $\lfloor \frac{1}{2(1+\delta)\gamma^2} \cdot \frac{m}{f} \rfloor > \ell$ . We have

$$D(S) < (1 + \epsilon)pn(S) \leq (1 + \epsilon)(1 + \delta) \cdot \frac{\gamma^2 f |S|}{T} \leq (1 + \epsilon) \cdot \frac{m}{2T}.$$

Since the epoch is assumed to be good,  $pn_u T \leq (1 + \delta)\gamma f$ ; also, the environment is  $(\gamma, s)$ -respecting, thus  $n(S) \leq \gamma n_u |S|$ . Putting these together verifies the second inequality, while the third follows from the bound on  $|S|$ .

With respect to the adversary, if  $\tau T B(J) < \epsilon m/4$ , then the total number of blocks is less than  $m$  and we are done. Otherwise, by Lemma 2(b),

$$B(J) < (1 + \epsilon)p|J| \leq (1 + \epsilon)(1 - \delta)pn(S) \leq (1 - \delta)(1 + \epsilon) \cdot \frac{m}{2T},$$

and the total count of blocks is again less than  $m$  by Condition C2.  $\square$

**Lemma 10.** *In a typical and hot-hand execution and a  $(\gamma, s)$ -respecting environment*

$$\text{GOODROUNDS}(r - 1) \implies \text{MEDIANTIME}(r).$$

*Proof.* Consider a chain  $\mathcal{C} \in \mathcal{S}_r$  with a target-recalculation point  $v \leq r$ . We prove that there is an honest median in  $\mathcal{C}$  computed in  $S = (v - \lfloor \epsilon^2 m/f \rfloor, v - \ell - 2\Delta)$ .

We will first argue that we may focus on the single target  $T$  of the epoch. Suppose—towards a contradiction—that there is a round  $w \in S$  in which an honest party drops  $\mathcal{C}$  to adopt a chain  $\mathcal{C}'$  of target  $T' \neq T$ . It follows that there is a target-recalculation point after  $\text{head}(\mathcal{C}' \cap \mathcal{C})$ . This cannot be on  $\mathcal{C}$ , because then the epoch would have lasted less than  $\lfloor \epsilon^2 m/f \rfloor$  rounds in violation of Lemma 9. Thus,  $\mathcal{C}'$  contains a target-recalculation point  $u \leq w$ . But then, no honest party would adopt  $\mathcal{C}$  during  $(u, v]$ . Since  $v - u > \ell + 2\Delta$ ,  $\mathcal{C}$  would become stale contradicting  $\mathcal{C} \in \mathcal{S}_r$ .

Since the execution is hot hand, there is a winning streak  $S^* \subseteq S$ . By Chain Growth the chain of each honest party increased by  $k_{\text{med}}$  blocks during  $S^*$ . Since  $S^*$  is a streak, no adversarial block lies between two of these blocks. We need to argue that all these blocks will belong to the chain of every honest party. If this is not the case, then it must be that either the first block of the streak or the last one belongs to a fork. In either case, suppose this block was computed in round  $w$ . By Lemma 6 there is an interval  $S$  such that  $w \in S$  and  $V(S) \leq 0$ . Let  $S' = S^* \cap S$  and  $S'' = S \setminus S'$ . Note that  $S'$  is nonempty since  $w \in S'$ . It follows from the definition of a winning streak that  $V(S') > 0$  and  $V(S'') \geq 0$ . Adding these inequalities we obtain a contradiction.  $\square$

**Lemma 11.** *In a typical and hot-hand execution and a  $(\gamma, s)$ -respecting environment*

$$\text{GOODROUNDS}(r-1) \implies \text{GOODCHAINS}(r).$$

*Proof.* Note that it is our assumption that the first round (the genesis) is a good target-recalculation point. Therefore, it suffices to show that if a recalculation point  $u$  in a chain  $\mathcal{C} \in \mathcal{S}_r$  is good, then the next one at  $v = u + \Lambda \leq r$  is also good. Let  $T$  be the target of the epoch starting at  $u$  and  $T'$  the target of the next one. We wish to show that  $f/2\gamma \leq pn_v T' \leq (1 + \delta)\gamma f$ . To that end, let  $u'$  and  $v'$  be the timestamps of the two target-recalculation points and set  $\Lambda' = v' - u'$ .

We prove first the lower bound. If  $\Lambda' \geq \gamma m/f$ , then  $T' \geq \gamma T$  (using  $\gamma \leq \tau$ ) and so  $pn_v T' \geq pn_u T'/\gamma \geq pn_u T \geq f/2\gamma$ , because  $u$  is assumed to be a good target-recalculation point. We assume now  $\Lambda' < \gamma m/f$ , which implies  $\Lambda' \leq (T'/T)(m/f)$ . Define  $S = [u, v]$ ,  $S' = [u - \ell - 2\Delta, v + \ell + 2\Delta]$ , and  $J$  the set of queries available to the adversary in  $S'$ . By Condition C1,

$$|S'| = \Lambda + 2\ell + 4\Delta \leq \Lambda + \frac{\epsilon^2 m}{16(1 + \delta)\gamma^2 f} \leq (1 + \epsilon^2/8)\Lambda,$$

where the last inequality follows from the lower bound on  $\Lambda$  implied by Lemma 9. Lemma 10 implies  $u' \geq u - \epsilon^2 m/f - \Phi$ , since  $u'$  will be forced by the honest median to be at most  $\epsilon^2 m/f$  rounds away from the timestamp of the median which is in turn at most  $\Phi$  rounds away from  $u$ . Also,  $v' \leq v + \Delta_{\text{fwd}}$ . Putting these together and using again the lower bound on  $\Lambda$  and Condition C1 as above

$$\frac{|\Lambda - \Lambda'|}{\Lambda'} \leq \frac{\epsilon^2 m/f + \Phi + \Delta_{\text{fwd}}}{m/2(1 + \delta)\gamma^2 f - \epsilon^2 m/f - \Phi - \Delta_{\text{fwd}}} \leq \frac{9\epsilon^2/8}{1/2(1 + \delta)\gamma^2 - 9\epsilon^2/8} < 6\epsilon^2,$$

where the last inequality uses Condition C2. By the last two displayed inequalities

$$|S'| \leq (1 + \epsilon^2/8)(1 + 6\epsilon^2)\Lambda' < (1 + \epsilon)\Lambda'.$$

Clearly, all blocks were computed during honest queries in  $S$  or adversarial ones in  $J$ . We now bound the contribution of each.

$$B(J) < (1 - \delta)(1 + \epsilon)pn(S') \leq (1 - \delta)(1 + \epsilon)p\gamma n_v |S'| \leq (1 - \delta)(1 + \epsilon)^2 p\gamma n_v \Lambda'.$$

Similarly,  $D(S) < (1 + \epsilon)pn(S) \leq (1 + \epsilon)p\gamma n_v \Lambda'$ . Assuming  $pn_v T' < f/2\gamma$  we obtain the contradiction

$$2\gamma pn_v \Lambda' \leq 2\gamma pn_v \cdot \frac{T'}{T} \cdot \frac{m}{f} < \frac{m}{T} \leq D(S) + B(J) < (2 + 2\epsilon - \delta)p\gamma n_v \Lambda'.$$

For the upper bound, let  $S = [u + \ell + 3\Delta, v - \ell + 3\Delta]$ . Note first that if  $\Lambda' \leq m/\gamma f$ , then  $T' \leq T/\gamma$  and so  $pn_v T' \leq p\gamma n_u T' \leq pn_u T \leq (1 + \delta)\gamma f$ , where we used that  $u$  is a good target-recalculation point. Thus, we may assume  $\Lambda' > m/\gamma f$ , which implies  $\Lambda' \geq (T'/T)(m/f)$ . Similarly to what we did above for the lower bound we may show  $|S| = |\Lambda - 2\ell - 6\Delta| \geq (1 - \epsilon)\Lambda'$ . Assuming  $pn_v T' > (1 + \delta)\gamma f$ , we obtain the contradiction

$$\frac{pn_v \Lambda}{(1 + \delta)\gamma} \geq \frac{pn_v}{(1 + \delta)\gamma} \cdot \frac{T'}{T} \cdot \frac{m}{f} > \frac{m}{T} \geq Q(S) > (1 - \epsilon)[1 - (1 + \delta)\gamma^2 f]^\Delta \cdot \frac{pn_v |S|}{\gamma}.$$

The first two inequalities have been discussed above. For the third one, note that since  $\mathcal{C} \in \mathcal{S}_r$ , by Lemma 7 there is a block computed by an honest party among the first and the last  $\ell + 2\Delta$  rounds of the epoch; the inequality follows by Chain Growth. The next one follows from Definition 7 ( $|S| \leq s$  due to Lemma 9) and so  $n(S) \geq n_v |S| / \gamma$ . The contradiction is a consequence of Condition C2 and the bound on  $|S|$ .  $\square$

**Lemma 12.** *In a typical and hot-hand execution and a  $(\gamma, s)$ -respecting environment*

$$\text{GOODROUNDS}(r-1) \implies \text{GOODROUNDS}(r).$$

*Proof.* Consider any  $\mathcal{C} \in \mathcal{S}_r$  and let  $u$  be its last recalculation point before  $r$  and  $T$  the associated target. Note that if  $r$  is a recalculation point, it follows directly by Lemma 11 that it is good. Otherwise, we need to show that  $f/2\gamma^2 \leq pn_r T \leq (1 + \delta)\gamma^2 f$ . By Lemma 11,  $f/2\gamma \leq pn_u T \leq (1 + \delta)\gamma f$ . By Lemma 9,  $n_u/\gamma \leq n_r \leq \gamma n_u$ . Combining these two bounds we obtain the desired inequality.  $\square$

**Theorem 3.** *Consider a typical and hot-hand execution in a  $(\gamma, 2(1+\delta)\gamma^2 m/f)$ -respecting environment. If the Conditions C1, C2, and C3 are satisfied, then all predicates of Definition 5 hold.*

*Proof.* We only need to verify that the predicates hold for the first  $\ell + 2\Delta$  rounds, assuming they hold at the first round. Note that if no epoch has been completed, all honest parties query for target  $T_0$  and are at most  $\gamma n_0$ . Thus, we only need to verify  $\text{DURATION}(\ell + 2\Delta)$ . The lower bound of Lemma 9 holds unless  $\text{GOODROUNDS}(r)$  fails for some  $r < \ell + 2\Delta$ , which does not happen by Lemma 12.  $\square$

#### 4.7 Common Prefix and Chain Quality

**Theorem 4 (Common Prefix).** *For a typical and hot-hand execution in a  $(\gamma, (1 + \delta)\gamma^2 m/f)$ -respecting environment, the common-prefix property holds for parameter  $\epsilon^2 m$ .*

*Proof.* Suppose common prefix fails for two chains  $\mathcal{C}_1$  and  $\mathcal{C}_2$  at rounds  $r_1 \leq r_2$ . It is not hard to see that in such a case there was a round  $r \leq r_2$  and two chains  $\mathcal{C}$  and  $\mathcal{C}'$  in  $\mathcal{S}_r$ , such that each had at least  $k$  blocks after  $\text{head}(\mathcal{C} \cap \mathcal{C}')$ . By Lemmas 9 and 8, at least  $\epsilon^2 m/2$  belong to one epoch. In view of Lemma 8, it suffices to show that that these were computed in at least  $\ell + 2\Delta$  rounds. Let  $T$  be the target of these blocks and suppose the honest parties query the oracle for target  $T$  during a set of rounds  $S$  of size  $\ell + 2\Delta$ . By Condition C1,  $|S| \leq \epsilon^2 m/16(1 + \delta)\gamma^2 f$ . Furthermore, by Theorem 3,  $pn_r T \leq (1 + \delta)\gamma^2 f$  holds for each  $r \in S$ . Putting these together, the number of such blocks that the honest parties computed are less than

$$T \cdot D(S) < (1 + \epsilon) \sum_{r \in S} pn_r T \leq (1 + \epsilon)\epsilon^2 m/16.$$

By Lemma 2 the adversary contributed less than  $(1 + \epsilon)\epsilon^2m/16(1 + \delta)$  blocks, for a total of less than  $\epsilon^2m/2$ .  $\square$

**Theorem 5 (Chain Quality).** *For a typical and hot-hand execution in a  $(\gamma, (1 + \delta)\gamma^2m/f)$ -respecting environment, the chain-quality property holds with parameters  $\ell + 2\Delta$  and  $\mu = \delta - 3\epsilon$ .*

*Proof.* Let us denote by  $B_i$  the  $i$ -th block of  $\mathcal{C}$  so that  $\mathcal{C} = B_1 \dots B_{\text{len}(\mathcal{C})}$  and consider  $K$  consecutive blocks  $B_u, \dots, B_v$ . Define  $K'$  as the least number of consecutive blocks  $B_{u'}, \dots, B_{v'}$  that include the  $K$  given ones (i.e.,  $u' \leq u$  and  $v \leq v'$ ) and have the properties (1) that the block  $B_{u'}$  was computed by an honest party at some round  $r$  or is  $B_1$  in case such block does not exist ( $r = 0$ ), and (2) that there exists a round  $r'$  such that  $B_1 \dots B_{v'} \in \mathcal{S}_{r'}$ . Denote by  $d'$  the total difficulty of these  $K'$  blocks. Define  $U = [r..r']$ ,  $S = [r + \Delta..r' - \Delta]$ , and  $J$  the adversarial queries in  $U$  starting with the first to obtain one of the  $K'$  blocks. Let  $x$  denote the total difficulty of all the blocks from honest parties that are included in the  $K$  blocks and—towards a contradiction—assume  $x < \mu d'$ . In a typical execution, all the  $K'$  blocks  $\{B_j : u' \leq j \leq v'\}$  have been computed in  $U$ . But then we have the following contradiction to Lemma 2(c).

$$A(J) \geq d' - x > (1 - \mu)d' \geq (1 - \mu)Q(S) = (1 - \delta + 3\epsilon)Q(S).$$

The first two inequalities follow from the definitions of  $x$  and  $d'$  and the assumed relation between them. It is not hard to see that the last inequality follows from Chain-Growth Lemma. Finally, to verify that this is indeed a contradiction, note that if  $U > (1 + \delta)\gamma^2m/f$  we may use Lemma 7 to partition  $U$  appropriately (using blocks computed by honest parties as pivot points) and apply Lemma 2(c) to each part. This is valid, since a block computed by an honest party provides both properties (1) and (2) required for  $K'$ .  $\square$

#### 4.8 Nakamoto Consensus: Consistency and Liveness

For parameters that satisfy Conditions C1 and C2, we can show that a typical and hot-hand execution in a  $(\gamma, (1 + \delta)\gamma^2m/f)$ -respecting environment enjoys Consistency and Liveness. Their proofs follow along the lines of previous work, such as [13, 24, 12]. In particular, Consistency will follow directly from the Common-Prefix property, that we show to hold in the above circumstances.

**Theorem 6 (Consistency).** *For a typical and hot-hand execution in a  $(\gamma, (1 + \delta)\gamma^2m/f)$ -respecting environment, Consistency is satisfied by setting the settled transactions to be those reported more than  $\epsilon m$  blocks deep.*

**Theorem 7 (Liveness).** *For a typical and hot-hand execution in a  $(\gamma, (1 + \delta)\gamma^2m/f)$ -respecting environment, Liveness is satisfied for depth  $\epsilon^2m$  with wait-time  $(4\gamma^2 + 1)\epsilon^2m/f$ .*

*Proof.* We claim that the chain  $\mathcal{C}$  of any honest party has at least  $\epsilon^2m$  blocks that were computed in the last  $4\epsilon^2\gamma^2m/(1 - 2\epsilon)f + 4\Delta$  rounds. Indeed,  $\mathcal{C}$  must



have a segment that lies in a single epoch—say of target  $T$ —and was computed in a set  $U$  of at least  $2\epsilon^2\gamma^2m/(1-2\epsilon)f+2\Delta$  consecutive rounds. If  $S$  is its subset without the first and last  $\Delta$  rounds, by Chain-Growth Lemma 1, the length of this segment is at least

$$T \cdot Q(S) > (1-\epsilon)[1-(1+\delta)f]^\Delta \sum_{r \in S} pn_r T \geq \frac{(1-2\epsilon)f|S|}{2\gamma^2} \geq \epsilon^2m.$$

Furthermore, if a transaction tx is included in any block computed by an honest party for the first  $\ell+2\Delta$  rounds, by Lemma 7, the chain  $\mathcal{C}$  of any honest party contains tx in a block  $B$ . The total wait-time amounts to

$$\ell + 6\Delta + \frac{4\epsilon^2\gamma^2m}{(1-2\epsilon)f} \leq \frac{\epsilon^2m}{2(1+\delta)f} + \frac{4\epsilon^2\gamma^2m}{(1-2\epsilon)f} \leq (4\gamma^2+1) \cdot \frac{\epsilon^2m}{f}. \quad \square$$

## References

1. L. Addario-Berry and B. A. Reed. *Ballot Theorems, Old and New*, pages 9–35. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
2. Adam Back. Hashcash. <http://www.cypherspace.org/hashcash>, 1997.
3. Christian Badertscher, Peter Gazi, Aggelos Kiayias, Alexander Russell, and Vassilis Zikas. Dynamic ad hoc clock synchronization. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology - EUROCRYPT 2021 - pages 399–428*. Springer, 2021.
4. Ran Canetti. Security and composition of multiparty cryptographic protocols. *J. Cryptology*, 13(1):143–202, 2000.
5. Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. Cryptology ePrint Archive, Report 2000/067, 2000. <http://eprint.iacr.org/2000/067>.
6. Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, pages 136–145*. IEEE Computer Society, 2001.
7. Amir Dembo, Sreeram Kannan, Ertem Nusret Tas, David Tse, Pramod Viswanath, Xuechao Wang, and Ofer Zeitouni. Everything is a race and nakamoto always wins. In Ligatti et al. [20], pages 859–878.
8. Devdatt Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, New York, 2009.
9. Cynthia Dwork, Nancy A. Lynch, and Larry J. Stockmeyer. Consensus in the presence of partial synchrony. *J. ACM*, 35(2):288–323, 1988.
10. Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In Ernest F. Brickell, editor, *CRYPTO*, volume 740 of *Lecture Notes in Computer Science*, pages 139–147. Springer, 1992.
11. Juan A. Garay and Aggelos Kiayias. Sok: A consensus taxonomy in the blockchain era. In Stanislaw Jarecki, editor, *Topics in Cryptology - CT-RSA 2020 - pages 284–318*. Springer, 2020.
12. Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 - pages 281–310*. Springer, 2015.

13. Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol with chains of variable difficulty. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - pages 291–323*. Springer, 2017.
14. Peter Gazi, Aggelos Kiayias, and Alexander Russell. Tight consistency bounds for bitcoin. In Ligatti et al. [20], pages 819–838.
15. Ari Juels and John G. Brainard. Client puzzles: A cryptographic countermeasure against connection depletion attacks. In *NDSS*. The Internet Society, 1999.
16. Jonathan Katz, Ueli Maurer, Björn Tackmann, and Vassilis Zikas. Universally composable synchronous computation. In Amit Sahai, editor, *Theory of Cryptography - 10th Theory of Cryptography Conference, TCC 2013, pages 477–498*. 2013.
17. Aggelos Kiayias and Giorgos Panagiotakos. Speed-security tradeoffs in blockchain protocols. *IACR Cryptology ePrint Archive*, 2015:1019, 2015.
18. Lucianna Kiffer, Rajmohan Rajaraman, and Abhi Shelat. A better method to analyze blockchain consistency. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS 2018, pages 729–744*. ACM, 2018.
19. Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.
20. Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors. *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9-13, 2020*. ACM, 2020.
21. Colin McDiarmid. *Probabilistic Methods for Algorithmic Discrete Mathematics*, chapter Concentration, pages 195–248. Springer 1998.
22. Michael Mitzenmacher and Eli Upfal. *Probability and computing - randomized algorithms and probabilistic analysis*. Cambridge University Press, 2005.
23. Satoshi Nakamoto. Bitcoin open source implementation of p2p currency. <http://p2pfoundation.ning.com/forum/topics/bitcoin-open-source>, 2009.
24. Rafael Pass, Lior Seeman, and Abhi Shelat. Analysis of the blockchain protocol in asynchronous networks. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology - EUROCRYPT 2017 - pages 643–673*, 2017.
25. Marshall C. Pease, Robert E. Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, 1980.
26. Ling Ren. Analysis of nakamoto consensus. *IACR Cryptology ePrint Archive*, 2019:943, 2019.
27. R. L. Rivest, A. Shamir, and D. A. Wagner. Time-lock puzzles and timed-release crypto. Technical report, Cambridge, MA, USA, 1996.
28. Fred B. Schneider. Implementing fault-tolerant services using the state machine approach: A tutorial. *ACM Comput. Surv.*, 22(4):299–319, 1990.
29. J.M. Steele and Mathematical Association of America. *The Cauchy-Schwarz Master Class: An Introduction to the Art of Mathematical Inequalities*. MAA problem books series. Cambridge University Press, 2004.
30. Lutz Warnke. On the method of typical bounded differences. *Combinatorics, Probability and Computing*, 25(2):269–299, 2016.

## A Mathematical Facts

All the following definitions and statements assume finite probability spaces and random variables with finite means.

**Definition 10.** [8, Definition 5.3] A sequence of random variables  $(X_0, X_1, \dots)$  is a martingale with respect to the sequence  $(Y_0, Y_1, \dots)$ , if, for all  $n \geq 0$ ,  $X_n$  is determined by  $Y_0, \dots, Y_n$  and  $\mathbf{E}[X_{n+1}|Y_0, \dots, Y_n] = X_n$ .

The following is closer to Theorem 3.15 in [21], but see also Theorems 8.1 and 8.2 in [8].

**Theorem 8.** Let  $(X_0, X_1, \dots)$  be a martingale with respect to the sequence  $(Y_0, Y_1, \dots)$ . Suppose an event  $G$  implies

$$X_k - X_{k-1} \leq b \text{ (for all } k) \quad \text{and} \quad V = \sum_k \mathbf{var}[X_k - X_{k-1}|Y_1, \dots, Y_{k-1}] \leq v,$$

Then, for non-negative  $n$  and  $t$ ,

$$\Pr[X_n \geq X_0 + t \wedge G] \leq \exp\left\{-\frac{t^2}{2v + 2bt/3}\right\}.$$

**Theorem 9.** [30, Theorem 1.6] Let  $X = (X_1, \dots, X_N)$  be a family of independent random variables with  $X_j$  taking values in a set  $\Lambda_j$  and let  $\Gamma = \prod_{j \in [N]} \Gamma_j$  where  $\Gamma_j \subseteq \Lambda_j$ . Assume there are numbers  $(c_j)_{j \in [N]}$  so that  $f : \prod_{j \in [N]} \Lambda_j \rightarrow \mathbb{R}$  satisfies the following. Whenever  $x, x' \in \prod_{j \in [N]} \Gamma_j$  differ only in the  $j$ -th coordinate and  $x, x' \in \Gamma$  we have  $|f(x) - f(x')| \leq c_j$  and  $|f(x) - f(x')| \leq d$  for all  $x, x' \in \prod_{j \in [N]} \Lambda_j$  that differ in at least one coordinate. Then, for all  $t \geq 0$ ,

$$\Pr[f(x) \leq \mathbf{E}[f(X)] - t - d\Pr[X \notin \Gamma]] \leq \exp\left\{-\frac{2t^2}{\sum_{j \in [N]} c_j^2}\right\} + \Pr[X \notin \Gamma].$$

The following Cauchy-Schwarz converse (see [29]) will be of use.

**Theorem 10.** For all non-negative real numbers  $a_k, b_k, k = 1, 2, \dots, n$  that satisfy  $m \leq a_k/b_k \leq M$  for some constants  $0 < m \leq M < \infty$ ,

$$\sqrt{\sum_{k=1}^n a_k^2} \sqrt{\sum_{k=1}^n b_k^2} \leq \frac{A}{G} \sum_{k=1}^n a_k b_k,$$

where  $A = (m + M)/2$  and  $G = \sqrt{mM}$ .

**Theorem 11 (Theorem 2 in [1]).** Suppose that  $X_1, \dots, X_n$  are integer-valued, independent, and identically distributed random variables with maximum value 1. For  $1 \leq i \leq n$ , let  $S_i = X_1 + \dots + X_i$ . Then, for any integer  $0 \leq k \leq n$ ,

$$\Pr[S_i > 0 \text{ for all } i = 1, 2, \dots, n | S_n = k] = \frac{k}{n}.$$

**Theorem 12 (Theorem 3 in [1]).** Let  $X_1, X_2, \dots$  be an infinite sequence of iid integer random variables with mean  $\mu > 0$  and maximum value 1 and for any  $i \geq 1$  let  $S_i = X_1 + \dots + X_i$ . Then

$$\Pr[S_i > 0 \text{ for } n = 1, 2, \dots] = \mu.$$

**Theorem 13.** Let  $a_1, \dots, a_n$  and  $b_1, \dots, b_n$  be nonnegative real numbers such that  $a_1 \geq \dots \geq a_n$  and

$$a_r + \dots + a_n \geq b_r + \dots + b_n \quad \text{for all } r = 1, \dots, n.$$

Let  $f(x)$  be an increasing concave function. Then

$$f(a_1) + \dots + f(a_n) \geq f(b_1) + \dots + f(b_n).$$

## B Proof of (All Executions are Typical) Theorem 1 and the Related Lemma 2

*Proof (of Theorem 1).* Since the length  $L$  of the execution is fixed we will prove the stated bound for a fixed set of consecutive rounds  $S$ —or, with respect to the adversary, a fixed set of consecutive queries  $J$ —and then apply a union bound over all such sets in the length of the execution. Furthermore, we may assume  $|S| \leq s$ . This is because  $\ell \leq s/2$  and we may partition  $S$  in parts such that each part has size between  $\ell$  and  $s$ . We then sum over all parts to obtain the desired bound. Let us also fix an execution  $E_0$  just before the beginning of  $S$  (or  $J$ ). We will prove that the statements fail with exponentially small probability for an arbitrary  $E_0$ . Note that  $E_0$  determines the number of parties  $n_0$  and  $t_0$  at the beginning of  $S$  (or  $J$ ) and the target  $T(J)$  associated with the first query in  $J$ .

For each round  $i \in S$ , define a Boolean random variable  $F_i$  equal to 1 exactly when all  $n_i$  hash values that were returned to the queries of the honest parties were above  $\min\{T : f(T, n_i) \geq (1 + \delta)\gamma^2 f\}$ ; define  $Z_i = Y_i \cdot F_{i+1} \cdots F_{i+\Delta-1}$ . Let  $G$  denote the event that the rounds in  $S$  are good. Given  $G$ , for any  $i \in S$ ,  $(F_i = 1) \implies (D_i = 0)$  and so  $Q_i \geq Z_i$ . Thus, for any  $d$ ,

$$\Pr\left[G \wedge \sum_{i \in [k]} Q_i \leq d\right] \leq \Pr\left[G \wedge \sum_{i \in [k]} Z_i \leq d\right],$$

and we may focus on the term on the right-hand side. Identify  $S$  with  $\{1, \dots, |S|\}$  and partition it with sets of the form  $S_j = \{j, j + \Delta, j + 2\Delta, \dots\}$  for  $j \in \{0, 1, \dots, \Delta - 1\}$ . We will show that, for each part  $S_j$ ,

$$\Pr\left[G \wedge \sum_{i \in S_j} Z_i \leq (1 - \epsilon)[1 - (1 + \delta)\gamma^2 f]^\Delta p \sum_{i \in S_j} n_i\right] \leq e^{-\lambda}.$$

Let us fix such a set  $S_j = \{s_1, s_2, \dots, s_\nu\}$ , with  $\nu \geq \lfloor |S|/\Delta \rfloor$ , and define the event  $G_t$  as the conjunction of the events  $G$  and  $t = \epsilon(1 - 2\gamma^2 f)^\Delta p n(S_j)$ . Note

that  $n(S_j) \leq L$  and so  $t$  ranges over a discrete set of size at most  $L$  and we can afford a union bound over it. Thus, it is sufficient to show that for any such  $t$ ,

$$\Pr\left[G_t \wedge \sum_{i \in S_j} Z_i \leq [1 - (1 + \delta)\gamma^2 f]^\Delta p \sum_{i \in S_j} n_i - t\right] \leq e^{-\lambda}.$$

To that end, consider the sequence of random variables

$$X_0 = 0; \quad X_k = \sum_{i \in [k]} Z_{s_i} - \sum_{i \in [k]} \mathbf{E}[Z_{s_i} | \mathcal{E}_{s_{i-1}}], \quad k \in [\nu].$$

This is a martingale with respect to the sequence  $\mathcal{E}_{s_1-1}(\mathcal{E}_0 = E_0), \dots, \mathcal{E}_{s_\nu-1}, \mathcal{E}$ , because (recalling basic properties of conditional expectation [21]),

$$\mathbf{E}[X_k | \mathcal{E}_{s_{k-1}}] = \mathbf{E}[Z_{s_k} - \mathbf{E}[Z_{s_k} | \mathcal{E}_{s_{k-1}}] | \mathcal{E}_{s_{k-1}}] + \mathbf{E}[X_{k-1} | \mathcal{E}_{s_{k-1}}] = X_{k-1}.$$

Specifically, the above follows from linearity of conditional expectation and the fact that  $X_{k-1}$  is a deterministic function of  $\mathcal{E}_{s_{k-1}+\Delta-1} = \mathcal{E}_{s_{k-1}}$ . Furthermore, given an execution  $E$  satisfying  $G_t$ ,

$$\epsilon \sum_{i \in S_j} \mathbf{E}[Z_i | \mathcal{E}_{s_{k-1}} = E_{s_{k-1}}] \geq \epsilon \sum_{i \in S_j} [1 - (1 + \delta)\gamma^2 f]^\Delta p n_i = t.$$

Thus, our goal is to show  $\Pr[-X_\nu \geq t \wedge G_t] \leq e^{-\lambda}$ .

We now provide the details relevant to Theorem 8. Consider an execution  $E$  satisfying  $G_t$  and let  $B$  denote the event  $\mathcal{E}_{s_{k-1}} = E_{s_{k-1}}$ . Note that  $Z_{s_k}^2 = Y_{s_k}^2 \cdot F_{s_k+1} \cdots F_{s_k+\Delta-1}$  and all these random variables are independent given  $B$ . Since  $X_k - X_{k-1} = Z_{s_k} - \mathbf{E}[Z_{s_k} | \mathcal{E}_{s_{k-1}}]$  and

$$Z_{s_k} - \mathbf{E}[Z_{s_k} | B] \leq \frac{1}{T_{s_k}^{\min}} = \frac{pn_{s_k}}{pn_{s_k} T_{s_k}^{\min}} \leq \frac{\gamma pn(S_j)}{pn_{s_k} T_{s_k}^{\min} |S_j|} \leq \frac{\gamma pn(S_j)}{\nu f / (2\gamma^2)} \leq \frac{2\gamma^3 t}{\epsilon(1 - 2\gamma^2)^\Delta f \nu} \stackrel{\text{def}}{=} b, \quad (6)$$

we see that the event  $G$  implies  $X_k - X_{k-1} \leq b$ . With respect to  $V = \sum_k \mathbf{var}[X_k - X_{k-1} | \mathcal{E}_{s_{k-1}}] \leq \sum_k \mathbf{E}[Z_{s_k}^2 | \mathcal{E}_{s_{k-1}}]$ , using the independence of the random variables and Proposition 1,

$$\sum_{k \in [\nu]} \mathbf{E}[Z_{s_k}^2 | B] \leq [1 - (1 + \delta)\gamma^2 f]^\Delta \sum_{k \in [\nu]} \frac{(pn_{s_k})^2}{pn_{s_k} T_{s_k}^{\min}} \leq \frac{[1 - (1 + \delta)\gamma^2 f]^\Delta}{f / (2\gamma^2)} \cdot \sum_{k \in [\nu]} (pn_{s_k})^2.$$

Applying Proposition 2 on this bound, we see that event  $G_t$  implies

$$V \leq \frac{2\gamma^3 [1 - (1 + \delta)\gamma^2 f]^\Delta}{f |S_j|} \cdot \left( \sum_{k \in [\nu]} pn_{s_k} \right)^2 \leq \frac{2\gamma^3 t^2}{\epsilon^2 f (1 - 2\gamma^2 f)^{\Delta+1} \nu} \stackrel{\text{def}}{=} v. \quad (7)$$

In view of these bounds (note that  $bt < \epsilon v$ ), by Theorem 8,

$$\Pr[-X_\nu \geq t \wedge G_t] \leq \exp\left\{-\frac{t^2}{2v(1 + \frac{\epsilon}{3})}\right\} \leq \exp\left\{-\frac{\epsilon^2 f [1 - (1 + \delta)\gamma^2 f]^\Delta \nu}{4\gamma^3 (1 + \frac{\epsilon}{3})}\right\} \leq e^{-\lambda},$$

where the last inequality follows from the value of  $\ell$  (recall that  $\nu \geq \ell/\Delta$  and equation (3)).

For the bound on  $D(S)$  it will be convenient to work per query. Let  $J$  denote the queries in  $S$ ,  $\nu = |J|$ , and  $Z_i$  the difficulty of any block obtained from query  $i \in J$ . Define the martingale sequence

$$X_0 = 0; \quad X_k = \sum_{i \in [k]} Z_i + \sum_{i \in [k]} \mathbf{E}[Z_i | \mathcal{E}_{i-1}], \quad k \in [\nu].$$

With similar calculations as above we obtain that  $G_t$  (with  $t = \epsilon p \nu$ ) implies

$$X_k - X_{k-1} \leq \frac{2\gamma^3 t}{\epsilon f |S|} \stackrel{\text{def}}{=} b \quad \text{and} \quad V \leq \frac{2\gamma^3 t^2}{\epsilon^2 f |S|} \stackrel{\text{def}}{=} v.$$

Applying Theorem 8 we obtain

$$\Pr[X_\nu \geq t \wedge G_t] \leq \exp\left\{-\frac{\epsilon t}{2b(1 + \frac{\epsilon}{3})}\right\} \leq e^{-\lambda}.$$

We next focus on part (b). For each  $j \in J$ , let  $A_j$  be equal to the difficulty of the block obtained with the  $j$ -th query as long as the target was at least  $T(J)/\tau$ ; thus,  $A(J) = \sum_{j \in J} A_j$ . If  $|J| = \nu$ , identify  $J$  with  $[\nu]$  and define the martingale

$$X_0 = 0; \quad X_k = \sum_{j \in [k]} A_j - \sum_{j \in [k]} \mathbf{E}[A_j | \mathcal{E}_{j-1}], \quad k \in [\nu].$$

For all  $k \in [\nu]$  we have  $X_k - X_{k-1} \leq \tau/T(J)$ ,  $\mathbf{var}[X_k - X_{k-1} | \mathcal{E}_{k-1}] \leq p\tau/T(J)$ , and  $\mathbf{E}[A_j | \mathcal{E}_{j-1}] \leq p$ . We may apply Theorem 8 with  $b = \tau/T(J)$ ,  $v = bp\nu \leq bt/\epsilon$ , and  $t = \max\{\epsilon p \nu, 2(\frac{1}{\epsilon} + \frac{1}{3})b\lambda\}$ . We obtain

$$\Pr\left[\sum_{j \in J} A_j \geq p\nu + t\right] \leq \exp\left\{-\frac{t}{2b(\frac{1}{3} + \frac{1}{\epsilon})}\right\} \leq e^{-\lambda}.$$

For part (c), as in [13], it can be shown that an insertion or a copy imply a collision, which can be shown to occur with probability at most  $(\frac{L}{2})2^{-\kappa}$ . Also, since there can be at most  $L$  predicted blocks, the probability a prediction occurs is at most  $L^2 2^{-\kappa}$ .  $\square$

*Proof (of Lemma 2).* (a) The middle inequality follows directly from the definition of the random variables. For the other two, let us assume first  $|U| \leq s$ . With respect to the lower bound on  $Q(S)$  we have

$$|J| < (1 - \delta)n(U) \leq (1 - \delta)\left(1 + \frac{2\gamma\Delta}{\ell}\right)n(S) < (1 - \delta)\left(1 + \frac{\epsilon^2}{2}\right)n(S).$$

The second inequality follows from  $|U| \leq s$  and Proposition 2; the third from  $\ell > 4\Delta\gamma^3/\epsilon^2 f$ , obtained from Equation (3). On the other hand, since  $|S| \geq \ell$

we may use Definition 7(a) to obtain  $Q(S) > (1 - 2\epsilon)pn(S)$ , which suffices for  $\epsilon \leq \delta/16$ .

With respect to the upper bound on  $D(U)$  we obtain similarly that

$$D(U) < (1 + \epsilon)pn(U) < (1 + \epsilon)\left(1 + \frac{\epsilon^2}{2}\right)pn(S) < (1 + 4\epsilon)Q(S).$$

For the case  $|U| > s$ , we partition the sets  $S$  and  $U$  into  $S_1, \dots, S_m$  and  $U_1, \dots, U_m$ , respectively, as follows. We consider any partition such that each part is at most  $s$  and at least  $\ell$  and  $S_i = U_i$  for  $i = 2, 3, \dots, m$  (this is always possible because  $s = \tau m/f > \ell/2$ , by Condition C1). The above derivations hold for each part and summing over all of them gives the desired inequalities.

(b) Either  $\epsilon p|J| \geq \tau\alpha(J)$  and Definition 7 applies directly, or  $p|J| < \tau\alpha(J)/\epsilon$  and by Equation (3) and Condition C1,

$$T(J) \cdot A(J) < \frac{2}{\epsilon^2} \cdot (1 + \epsilon)\left(1 + \frac{\epsilon}{3}\right)\tau\lambda < \frac{f\ell}{2\gamma^3} \leq \frac{(1 - \epsilon)^3}{32(1 + \delta)^2\gamma^9} \cdot \epsilon^2 m.$$

(c) In a  $(\gamma, s)$ -respecting environment,  $\gamma n(S) \geq n_w|S|$ . Incorporating this in Definition 7 we obtain the first bound. For the second one, Using  $\epsilon < 1/6$ ,

$$\epsilon(1 - 2\epsilon)pn(S) \geq \epsilon(1 - 2\epsilon)\frac{pn_w|S|}{\gamma} \geq \frac{\epsilon(1 - 2\epsilon)f\ell}{2\gamma^3 T(J)} > \tau\alpha(J).$$

As in (a),  $p|J| \leq (1 - \delta + \epsilon^2/2)pn(S)$ . We obtain  $A(J) \leq (1 - \delta + \epsilon)pn(S)$  and use Condition C2.

(d) Suppose the parties query the oracle for target  $T$  during a set of rounds  $S$  of size  $\ell + 2\Delta$  and so we may bound its size by the right-hand side of Condition C1. Furthermore, since the blocks belong to a good chain,  $pn_r T \leq (1 + \delta)\gamma^2 f$  for each  $r \in S$ . Putting these together, the number of such blocks that the honest parties computed are less than

$$T \cdot D(S) < (1 + \epsilon) \sum_{r \in S} pn_r T \leq (1 + \epsilon) \cdot \frac{(1 - \epsilon)^3}{16(1 + \delta)\gamma^4} \cdot \epsilon^2 m \leq \frac{(1 - \epsilon)^2}{16(1 + \delta)\gamma^4} \cdot \epsilon^2 m.$$

Adding the contribution of the adversary and using the bound in part (b), we obtain the desired bound.  $\square$

## C The Bitcoin Backbone Protocol with Variable Difficulty (cont'd)

In this section we give a more detailed description of the Bitcoin backbone protocol with chains of variable difficulty. The presentation is based on the description in [12].

## C.1 The protocol

As in [12] in our description of the backbone protocol we intentionally avoid specifying the type of values/content that parties try to insert in the chain, the type of chain validation they perform (beyond checking for its structural properties with respect to the hash functions  $G(\cdot), H(\cdot)$ ), and the way they interpret the chain. These checks and operations are handled by the external functions  $V(\cdot), I(\cdot)$  and  $R(\cdot)$  (the *content validation function*, the *input contribution function* and the *chain reading function*, resp.) which are specified by the application that runs “on top” of the backbone protocol.

The Bitcoin backbone protocol in the dynamic setting is specified as Algorithm 4 and depends on three sub-procedures.

**Chain validation.** The `validate` algorithm performs a validation of the structural properties of a given chain  $\mathcal{C}$ . It is given as input the value  $q$ , as well as hash functions  $H(\cdot), G(\cdot)$ . It is parameterized by the content validation predicate  $V(\cdot)$  as well as by  $D(\cdot)$ , the *target calculation function* (see Section 3). For each block of the chain, the algorithm checks that the proof of work is properly solved (with a target that is suitable as determined by the target calculation function), and that the counter  $ctr$  does not exceed  $q$ . Furthermore it collects the inputs from all blocks,  $\mathbf{x}_{\mathcal{C}}$ , and tests them via the predicate  $V(\mathbf{x}_{\mathcal{C}})$ ; note that  $V(\varepsilon) = \mathbf{true}$ . Chains that fail these validation procedure are rejected. (Algorithm 1.)



---

**Algorithm 1** The *chain validation predicate*, parameterized by  $g, D$ , the hash functions  $G(\cdot), H(\cdot)$ , and the *input validation predicate*  $V(\cdot)$ . The input is chain  $\mathcal{C}$ .

---

```

1: function validate( $r_{\text{now}}, \mathcal{C}$ )
2:    $valid \leftarrow V(\mathbf{x}_{\mathcal{C}})$ 
3:   if  $valid = \text{True} \wedge (\mathcal{C} \neq \varepsilon)$  then  $\triangleright \mathcal{C}$  is non-empty and meaningful w.r.t.  $V(\cdot)$ 
4:      $r_{\text{R}} \leftarrow r_{\text{now}} + \Delta_{\text{fwd}}$ 
5:      $r_{\text{L}} \leftarrow 0$ 
6:      $len \leftarrow |\mathcal{C}|$ 
7:      $\langle r', st', x', ctr' \rangle \leftarrow \langle 0, \perp, \text{"genesis"}, \perp \rangle$ 
8:      $i \leftarrow 1$ 
9:      $T \leftarrow T_{\text{initial}}$ 
10:    while  $(i \leq len) \wedge valid$  do
11:       $\langle r, st, x, ctr \rangle \leftarrow \text{head}(\mathcal{C}^{\lceil len-i \rceil})$   $\triangleright$  Get the  $i$ -th block; note  $\mathcal{C}^{\lceil 0 \rceil} = \mathcal{C}$ 
12:       $st_{\text{prev}} \leftarrow H(ctr', G(r', st', x'))$   $\triangleright$  Calculate the hash of previous block
13:      if  $\text{validblock}_q^T(\langle r, st, x, ctr \rangle) \wedge (st_{\text{prev}} = st) \wedge (r_{\text{L}} < r < r_{\text{R}})$  then
14:         $\langle r', st', x', ctr' \rangle \leftarrow \langle r, st, x, ctr \rangle$   $\triangleright$  Retain current block
15:         $timeseq \leftarrow \langle r \rangle || timeseq$   $\triangleright$  Prepend timestamp
16:        if  $|timeseq| > k_{\text{med}}$  then  $\triangleright$  We have enough timestamps for median
17:           $timeseq \leftarrow timeseq^{\downarrow 1}$   $\triangleright$  Remove the last element
18:        end if
19:         $r_{\text{L}} \leftarrow \max\{r_{\text{L}}, \text{median}(timeseq) + 1\}$   $\triangleright$  Advance left time bound
20:      else
21:         $valid \leftarrow \text{False}$   $\triangleright$  Blockchain is not valid
22:      end if
23:       $T \leftarrow D(\mathbf{r}_{\mathcal{C}^{\lceil len-i \rceil}})$   $\triangleright$  Calculate next target
24:       $i \leftarrow i + 1$ 
25:    end while
26:  end if
27:  return  $valid$ 
28: end function

```

---

Note that in the Bitcoin implementation it holds that  $\Delta_{\text{fwd}} = 2$  hours and  $k_{\text{med}} = 11$ .

**Chain comparison.** The objective of the second algorithm, called `maxvalid`, is to find the “best possible” chain when given a set of chains. The algorithm is straightforward and is parameterized by a `max(·)` function that applies some ordering in the space of chains. The most important aspect is the chains’ difficulty in which case `max(C1, C2)` will return the most *difficult* of the two. In case `diff(C1) = diff(C2)`, some other characteristic can be used to break the tie. In our case, `max(·, ·)` will always return the first operand to reflect the fact that parties adopt the first chain they obtain from the network. (Algorithm 2.)

---

**Algorithm 2** The function that finds the “best” chain, parameterized by function  $\max(\cdot)$ . The input is  $\{\mathcal{C}_1, \dots, \mathcal{C}_k\}$ .

---

```

1: function maxvalid( $r_{\text{now}}, \mathcal{C}_1, \dots, \mathcal{C}_k$ )
2:    $temp \leftarrow \varepsilon$ 
3:   for  $i = 1$  to  $k$  do
4:     if validate( $r_{\text{now}}, \mathcal{C}_i$ ) then
5:        $temp \leftarrow \max(\mathcal{C}, temp)$ 
6:     end if
7:   end for
8:   return  $temp$ 
9: end function

```

---

**Proof of work.** The third algorithm, called **pow**, is the proof of work-finding procedure. It takes as input a chain and attempts to extend it via solving a proof of work. This algorithm is parameterized by two hash functions  $H(\cdot), G(\cdot)$  as well as the parameter  $q$ . Moreover, the algorithm calls the target calculation function  $D(\cdot)$  in order to determine the value  $T$  that will be used for the proof of work. The procedure, given a chain  $\mathcal{C}$  and a value  $x$  to be inserted in the chain, hashes these values to obtain  $h$  and initializes a counter  $ctr$ . Subsequently, it increments  $ctr$  and checks to see whether  $H(ctr, h) < T$ ; in case a suitable  $ctr$  is found then the algorithm succeeds in solving the POW and extends chain  $\mathcal{C}$  by one block. (Algorithm 3.)

---

**Algorithm 3** *Proof of work* single step based on hash functions  $H(\cdot), G(\cdot)$  and target calculation function  $D(\cdot)$ .

---

```

1: function pow( $r_{\text{now}}, x, \mathcal{C}$ )
2:   if  $\mathcal{C} = \varepsilon$  then                                     ▷ Determine proof of work instance.
3:      $prev \leftarrow 0$ 
4:      $ctr \leftarrow 0$ 
5:   else
6:      $\langle r', prev', x', ctr' \rangle \leftarrow \text{head}(\mathcal{C})$ 
7:      $prev \leftarrow H(ctr', G(r', prev', x'))$ 
8:   end if
9:    $B \leftarrow \varepsilon$ 
10:   $T \leftarrow D(r_c)$                                      ▷ Calculate target for next block based on timestamps.
11:   $h \leftarrow G(r, prev, x)$ 
12:  if  $(H(ctr, h) < T)$  then                               ▷ This  $H(\cdot)$  invocation is subject to 1 query/round.
13:     $B \leftarrow \langle r, prev, x, ctr \rangle$ 
14:  end if
15:   $ctr \leftarrow ctr + 1 \pmod{2^{32}}$ 
16:   $\mathcal{C} \leftarrow \mathcal{C}B$                                      ▷ Chain is extended, if  $B \neq \varepsilon$ 
17:  return  $\mathcal{C}$ 
18: end function

```

---

**The backbone protocol.** The core of the protocol is similar to that of [13]. We recall some of the main functions as well as point to the new elements we have added compared to previous abstractions. Parties always check the **active** flag to make sure they detect they have missed one or more rounds. In case the **active** flag is false, they broadcast a special message ‘**Join**’ (that requests the most recent version of the blockchain from other other parties) and enter into bootstrapping mode which lasts for a certain period of time denoted by  $\Delta_{\text{bootstr}}$ . To respond to such request, when online parties receive the special request message in their RECEIVE() tape they broadcast their chain. The input contribution function  $I(\cdot)$  and the chain reading function  $R(\cdot)$  are applied to the values stored in the chain; we are not concerned with the way these functions are defined.

As in past work, [12, 13] the input tape of a party contains two types of symbols, READ and (INSERT, *value*); other inputs are ignored. A READ results to the party applying function  $R(\cdot)$  to its current chain and writing the result onto the output tape OUTPUT(), while a (INSERT, *value*) symbol is taken into account by the  $I(\cdot)$  function when it determines the contribution of the party in extending its chain.

One of the novel elements in the current treatment is the fact that the current time is always determined by querying the clock functionality, which allows the adversary to apply a drift on a party’s clock up to the  $\Phi$  bound.

The pseudocode of the backbone protocol is presented in Algorithm 4.

---

**Algorithm 4** The Bitcoin backbone protocol in the dynamic setting at round “round” on local state  $(st, \mathcal{C})$  parameterized by the *input contribution function*  $I(\cdot)$  and the *chain reading function*  $R(\cdot)$ . The “active” flag is False if and only if the party was inactive in the previous round. **bootstrapping** is initially False.

---

```

1:  $currenttime \leftarrow RequestTime$  ▷ request time from the clock functionality
2:  $bootstrapping \leftarrow bootstrapping \wedge active$ 
3: if  $active = True \wedge \neg bootstrapping$  then ▷ Ledger maintenance mode
4:    $DIFFUSE(Ready)$ 
5:    $round \leftarrow currenttime$ 
6:    $\tilde{\mathcal{C}} \leftarrow \maxvalid(round, \mathcal{C}, \text{all chains } \mathcal{C}' \text{ found in } RECEIVE())$ 
7:    $round \leftarrow \max\{round, 1 + \text{median of } \tilde{\mathcal{C}} \text{ last } 2k_{med} - 1 \text{ blocks}\}$ 
8:   if  $INPUT()$  contains READ then
9:      $write R(x_c)$  to  $OUTPUT()$ 
10:  end if
11:   $\langle st, x \rangle \leftarrow I(st, \tilde{\mathcal{C}}, round, INPUT(), RECEIVE())$ 
12:   $\mathcal{C}_{new} \leftarrow pow(round, x, \tilde{\mathcal{C}})$ 
13:  if  $(\mathcal{C} \neq \mathcal{C}_{new}) \vee (Join \in RECEIVE())$  then
14:     $\mathcal{C} \leftarrow \mathcal{C}_{new}$ 
15:     $DIFFUSE(\mathcal{C})$  ▷ chain is diffused when updated or when someone joins.
16:  end if
17:   $DIFFUSE(RoundComplete)$ 
18: else ▷ Bootstrapping mode
19:    $active \leftarrow True$ 
20:   if  $bootstrapping = True$  then ▷ Node is in the process of bootstrapping
21:      $timelapsed \leftarrow timelapsed + \max\{currenttime - timelapsed, 0\}$ 
22:   else ▷ Node just woke up and needs to bootstrap
23:      $bootstrapping \leftarrow True$ 
24:      $timelapsed \leftarrow 0$ 
25:   end if
26:   if  $timelapsed > \Delta_{bootstr}$  then
27:      $bootstrapping \leftarrow False$  ▷ Node is ready to engage
28:   else
29:      $DIFFUSE(Join, RoundComplete)$  ▷ Node is asking for blocks to synchronize
30:   end if
31: end if

```

---