Permissionless Clock Synchronization with Public Setup

Juan Garay^{1*}, Aggelos Kiayias², and Yu Shen^{3**}

¹ Texas A&M University garay@cse.tamu.edu
 ² University of Edinburgh and IOHK aggelos.kiayias@ed.ac.uk
 ³ University of Edinburgh yu.shen@ed.ac.uk

Abstract. The permissionless clock synchronization problem asks how it is possible for a population of parties to maintain a system-wide synchronized clock, while their participation rate fluctuates —possibly very widely— over time. The underlying assumption is that parties experience the passage of time with roughly the same speed, but however they may disengage and engage with the protocol following arbitrary (and even chosen adversarially) participation patterns. This (classical) problem has received renewed attention due to the advent of blockchain protocols, and recently it has been solved in the setting of proof of stake, i.e., when parties are assumed to have access to a trusted PKI setup [Badertscher *et al.*, Eurocrypt '21].

In this work, we present the first proof-of-work (PoW)-based permissionless clock synchronization protocol. Our construction assumes a public setup (e.g., a CRS) and relies on an honest majority of computational power that, for the first time, is described in a fine-grain timing model that does not utilize a global clock that exports the current time to all parties. As a secondary result of independent interest, our protocol gives rise to the first PoW-based ledger consensus protocol that does not rely on an external clock for the time-stamping of transactions and adjustment of the PoW difficulty.

1 Introduction

In the classical clock synchronization problem, thoroughly studied over the past four decades by the distributed computing community—non-exhaustively, [18,19,16,8,25,26,1,20]—, a set of processors, each one possessing a timer that is within a bounded rate of drift from "nominal time" (the real time—called Newtonian time in [8]), should realize logical clocks that are within a distance Skew $\in \mathbb{N}$ of each other and within a linear envelope of nominal time. The typical threat model involves a subset of parties who deviate arbitrarily either from correct protocol execution or in terms of their clock speed and may as a

 $^{^{\}star}$ Research supported by NSF grants no. 2001082 and 2055694.

^{**} Work supported by Input Output – IOHK through their funding of the Edinburgh Blockchain Technology Lab.

result prevent synchronization from happening. A clock synchronization protocol has parties exchanging messages to suitably adjust their clocks so that the synchronization condition is achieved.

Up until the work of [3] all prior work in clock synchronization assumed that the number of parties are known during the protocol execution (and available, unless they are assumed adversarial⁴). This standard assumption in Byzantine fault tolerance protocols was challenged first with the advent of the Bitcoin blockchain and related "permissionless" protocols. As exemplified in [11,12], the Bitcoin blockchain operates in a setting where the number of active parties may be unknown and continuously fluctuating throughout the protocol execution. While such results paved the way to rethink the problem of consensus in this setting (cf. [24,15]), near perfectly synchronized clocks remained a central assumption in all previous security analyses of blockchain protocols (cf. [10,4,22,11,12]).

In the setting where participation is dynamic and fluctuating over time, the adversary can introduce and remove honest parties at will without notifying the existing participants. As a result, existing clock synchronization algorithms (e.g., [25,1,20]) do not directly translate to such permissionless setting because they fundamentally rely on the fact that the parties are aware of the number of parties as well as of the number of tolerated corruptions/faults—i.e., they are able to *count*—and a different protocol design technique is needed.

The main challenge in this transient participation setting shifts from correcting the bounded-rate drift occurring between the ever connected honest parties over time to the task of bringing up to sync freshly joining parties who start without any information about nominal time, while accommodating for the fact that a (possibly large) number of honest parties is no longer active. In [3], assuming a so-called private-state setup [15] (specifically, a PKI), a protocol called "Ouroboros Chronos" is presented that can synthesize a notion of global time using a continuous flow of clock measurements that are provided by parties who only transiently participate in the protocol and their local clocks are assumed to be correct up to a bound. The level of participation fluctuates broadly with the only requirements that (i) it does not become negligible, and (ii) honest majority is preserved in terms of stake (all parties have a number of coins associated to their public keys that amount to their individual stake). Given this, their result leaves open the question of only utilizing a public(-state) setup.

To our knowledge, the only known result with a public setup in the permissionless setting, again from [3], is that parties may use a Nakamoto-style longest chain blockchain without difficulty adjustment and use the block index to define a concept of global time. The obvious downside of this idea is that the protocol execution speeds up and slows down as participation fluctuates and,

⁴ We note that the problem of joining parties in the context of clock synchronization was considered, but only conditionally on the new party agreed upon and approved by a sufficient number of participants; see [16].

most importantly, it will be entirely insecure when there is a steady increase (or decrease) of participants, making the construction essentially only suitable for a static model where the number of parties (i.e., the computational power invested in the system for proof of work) remains fixed.

This motivates the current work, where the following open question is being tackled:

Is it possible for a dynamically changing population of peers to synchronize their clocks utilizing only a public setup and assuming PoW?

One apparent difficulty in answering this question is that using a blockchain protocol to derive consistency for clock adjustments runs into the complication that the blockchain protocol itself utilizes a clock to adjust the PoW difficulty at regular intervals. Indeed, the Bitcoin blockchain [21] relies fundamentally on a global clock being available to all parties.⁵ It follows that this observation suggests also a secondary open question that will be tackled as well:

Is there a blockchain protocol in the PoW setting that has no dependency on a publicly accessible global clock?

1.1 Overview of Our Results

The clock synchronization problem asks parties to report clocks that satisfy two properties (cf. [8]) (i) *bounded skew:* the parties maintain logical clocks whose difference is upper bounded, and (ii) *linear envelope synchronization:* the logical clock reported by a party is always within a *linear envelope* of the nominal time. Note that we are interested in a formulation of this problem in a very general setting where some parties are adversarial and hence deviate from the protocol arbitrarily, while honest parties may come and go following arbitrary participation patterns. Given this setting we formulate the *desideratum* of a synchronized clock only with respect to a class of parties we call *alert*, which are honest parties that have also been online for a sufficiently long time to catch up with all protocol messages. More formally, the clock synchronization problem is stated as follows.

Definition 1 (Clock Synchronization). There exist constants Skew $\in \mathbb{N}$, shiftLB, shiftUB $\in (0, 1)$ such that honest parties' logical clocks satisfy the following two properties:

- Bounded skews. Let $\mathbf{r}_1, \mathbf{r}_2$ be the reported logical clocks of two alert parties at any nominal time r. Then $|\mathbf{r}_1 - \mathbf{r}_2| \leq Skew$.

⁵ The protocol implements such clock by having nodes querying other nodes in the network and possibly seeking user input — it has no way of deriving a clock from the protocol operation itself. See [12] for more details.

- 4 Juan Garay, Aggelos Kiayias, and Yu Shen
- Linear envelope synchronization. Each alert party's logical clock stays in a (U, L)-linear envelope⁶ with respect to the nominal time r, where U = 1/(1 - shiftUB) and L = 1/(1 + shiftLB).

Solving the clock synchronization problem asks for a protocol that within a certain threat model achieves the two properties. This brings us to our first main result.

A model for permissionless clock synchronization in the PoW setting. Our model (Section 2) simultaneously facilitates (i) the dynamic participation of parties, (ii) imperfect local clocks, and (iii) resource bounding by restricting parties' queries to a random oracle (cf. [13]). Specifically, we extend the previous model of the global imperfect clock of [3] to the PoW setting by introducing a random oracle functionality that apportions random oracle queries per unit of time between the honest parties and the adversary in a manner consistent with an honest majority assumption in terms of computational power. The concept of time provided by the imperfect local clock functionality of [3] enables parties to advance their local clocks and experience time at roughly the same speed (a maximum drift of $\Phi_{\rm clock}$ is allowed). Note that the environment is allowed to introduce new parties and remove old parties at will, something that results in them being de-registered from the clock functionality; when this happens the clock functionality is not responsible for keeping them up to speed with the rest of the honest parties. In this way, parties can be seen as entirely transiently engaging with the protocoleach individual party may only engage for a small fraction of the total execution time as adaptively decided by the environment. Armed with our model, we then present our second main result.

A new protocol for permissionless clock synchronization in the PoW model. We describe our new PoW-based clock synchronization protocol Timekeeper in Section 3. The construction is based on three key ingredients: (i) A mechanism that repurposes the concept of 2-for-1 PoWs introduced in [10] and subsequently used to achieve various properties such as fairness in [23] and high throughput in [5], to the setting of time-keeping by employing it to enable the collection of "timing beacons" from the active parties in a rolling window process; (ii) a PoW-based longest-chain type of blockchain that enables parties at regular intervals to reach consensus about the timing beacons that are shared and extract a suitable correction to their local clocks taking into account the arrivals of the beacons; and (iii) a novel target-recalculation function that can be thought of as the *reverse* of the one used in Bitcoin, that uses protocol recorded timestamps as a means to define the length of an epoch, and then uses the number of blocks produced in that period of time to adjust the PoW difficulty accordingly.

Putting these elements together, our clock synchronization protocol instructs parties when their local clock passes some specific moment (which happens pe-

⁶ A function $f : \mathbb{R} \to \mathbb{R}$ is within a (U, L)-linear envelope if and only if it holds that $L \cdot x \leq f(x) \leq U \cdot x$, for all x.

riodically with respect to the interval length) to execute an adjustment on their local clock based on the median value of the beacon timestamps and their corresponding arrival time. Moreover, towards the goal of letting newly joining parties become synchronized with the protocol time, we present a joining procedure which requires the fresh parties to passively listen to the protocol execution for a while and then synchronize with other honest participants.

Based on the ledger consensus function offered by our protocol it is easy to derive also the following result.

A new PoW-based blockchain protocol without a global clock. Given that our protocol is a Nakamoto-style "most-difficult chain" type of protocol that faclitates clock synchronization, it is easy to transform it to a full-fledged blockchain protocol that admits transactions as in Bitcoin script or Ethereum smart contracts. The resulting blockchain has the novel property that it does not depend on accessing a globally available clock. Instead, parties utilize their local clocks which may be drifting or be out of sync, but thanks to the synchronization (sub-)protocol that is offered by our construction they can adjust their local clocks periodically. This eliminates time as an attack vector in the context of PoWbased blockchain protocols and demonstrates that it is possible to achieve ledger consensus using merely local clocks in a fully dynamic setting where parties may come and go adaptively per the adversary's instructions.

Security analysis. We present the full security analysis of Timekeeper in Section 4. As a high-level overview, we proceed to adapt the analytical toolset from [11,12] to the imperfect-local-clock model. Notably, we modify the concept of target recalculation epoch boundaries (from "point" to "zone") and the concept of isolated successes (which addresses the question of under what circumstances can a hash success guarantee the increase of accumulated difficulty). As an intermediate step, we study several predicates aiming at providing the "good" properties of an execution starting from the onset and until a given nominal round.

Our inductive-style proof works in the following manner. We prove that if at the onset, the PoW difficulty is appropriately set and the steady block-generation rate lasts during the whole clock synchronization interval, parties can maintain good skews after they enter the next interval and the shift value they compute to adjust their clocks is properly bounded. In addition, if good skews and certain time adjustment calculations are maintained during a target recalculation epoch, the block production rate will be properly controlled in the next epoch. To sum up, this guarantees that "good" properties can be achieved during the whole execution given a "safe" start and a bounded change in the number of parties (which can nevertheless still be exponential). We also provide an analysis of the joining procedure showing that joining parties starting with no *a-priori* knowledge of global time, can listen in and bootstrap their logical clock, turning themselves into *alert* parties being capable of fully engaging with the protocol.

In summary, **Timekeeper** solves the clock synchronization problem as defined above as follows.

Theorem (Theorem 3, informal). Let Φ_{clock} be the maximum drift allowed on parties' local clocks and Δ the maximum (and unknown) message transmission delay. Then Timekeeper solves the clock synchronization problem assuming bounded dynamic participation and an honest majority in terms of random oracle queries, with parameter values

Skew = 2Φ , shiftLB = $3\Phi/R$, shiftUB = $2\Phi/R$,

where $\Phi = \Delta + \Phi_{clock}$ and $R \in \mathbb{N}^+$ is a parameter chosen sufficiently large w.r.t. the security parameter and reflects the time required for an honest party to become alert.

Organization of the paper. The rest of the paper is organized as follows. In Section 2 we present our model, relevant definitions and building blocks. We describe our Timekeeper protocol in Section 3 and present the full analysis in Section 4. Detailed description of protocols, functionalities, and proofs can be found in the full version of the paper [14].

2 Model and Building Blocks

In this paper we adapt the timing and networking model of [3] to the setting of proof of work, obviating the requirement for a PKI as a setup assumption. In more detail, in the model there is an upper bound Δ in message transmission (cf. [9,22,4,12]), and parties do not have access to a global clock, but instead rely on their local clocks, whose drift is assumed to be upper-bounded by Φ_{clock} . What complicates matters is that the model supports dynamic participation where parties may join and leave during the protocol execution without warning (it is worth noting here that this is where the difficulty of our setting is derived from: indeed if all honest parties were online throughout then it would be trivial to implement a logical clock by incrementing a counter). For succinctness, we choose to express primitives and building blocks (see below) in our execution model utilizing the ideal functionality language of [7], but we do not pursue a composability analysis for our security properties, which are expressed in a game based manner as in [10,22].

2.1 Imperfect Local Clocks

As in [3], and as mentioned above, in this paper we remove the assumption that parties have access to a global clock, as in [10,11,22,4,12], and instead assume *imperfect local clocks*. In a nutshell, every honest party maintains a local clock variable by communicating with an imperfect local clock functionality $\mathcal{F}_{\text{ILCLOCK}}$. In contrast to the global-setup clock functionality in [17], where parties learn the

exact global time and thus strong synchrony is guaranteed, parties registered with $\mathcal{F}_{\text{ILCLOCK}}$ will only receive "ticks" from the functionality to indicate that they should update their own clocks. In addition, $\mathcal{F}_{\text{ILCLOCK}}$ issues "imperfect" ticks, i.e., the adversary is allowed to set a bounded drift to each party by manipulating its corresponding status variable in $\mathcal{F}_{\text{ILCLOCK}}$. $\mathcal{F}_{\text{ILCLOCK}}$ can be viewed as a variant from [3]'s with adaptations to provide a more natural clock model with real-word resources and in the proof-of-work setting.

For a detailed description of the functionality, see [14]. Here we just elaborate on the "imperfect" aspect of the clock and on the adversarial manipulation of clock drifts. Specifically, we allow the adversary to set some drifts to parties' local clocks, which will accelerate or stall their progress; such values are globally bounded by Φ_{clock} . This assumption allows local clocks to proceed at "roughly" the same speed.

Further, the adversary \mathcal{A} can adaptively manipulate the drift of honest parties' clocks by sending CLOCK-FORWARD and CLOCK-BACKWARD messages to the functionality⁷ after they conclude the current round. If \mathcal{A} issues CLOCK-FORWARD for party P, it will enter a new local round before $\mathcal{F}_{\text{ILCLOCK}}$ updates the nominal time, and this can be repeated as long as P's drift is not Φ_{clock} rounds larger than other honest parties. On the other hand, if \mathcal{A} issues CLOCK-BACKWARD, it will set P's budget to a negative value, thus preventing $\mathcal{F}_{\text{ILCLOCK}}$ from updating d_{P} at the end of the nominal round (d_{P} is the functionality variable that captures whether the party P has made its move for the clock tick). I.e., P will still be in the same logical round during these two nominal rounds. Again, this process can be repeated by \mathcal{A} as long as the drift on P is not Φ_{clock} rounds smaller than others. As a consequence, the targeted party's local clock may remain static for several nominal rounds.

2.2 Other Core Functionalities

Common Reference String. We model a public-state setup by the CRS functionality $\mathcal{F}_{CRS}^{\mathcal{D}}$. The functionality is parameterized with some distribution \mathcal{D} with sufficiently high entropy. Once $\mathcal{F}_{CRS}^{\mathcal{D}}$ receives (RETRIEVE, sid) from either the adversary \mathcal{A} or a party P for the first time, it generates a string $d \leftarrow \mathcal{D}$ as the common reference string. In addition, $\mathcal{F}_{CRS}^{\mathcal{D}}$ will immediately send a message (RETRIEVED, sid) to functionality $\mathcal{W}(\mathcal{F}_{RO})$ (described next) to indicate that $\mathcal{W}(\mathcal{F}_{RO})$ should start to limit the adversarial RO queries. For all subsequent activations, $\mathcal{F}_{CRS}^{\mathcal{D}}$ simply returns d to the requester.

(Wrapped) Random Oracle. By convention, we model parties' calls to the hash function used to generated proofs of work as assuming access to a random oracle; this is captured by the functionality \mathcal{F}_{RO} . Notice that with regards to bounding

⁷ As such, our clock functionality is a more natural model of the real world compared to [3]'s, as it allows \mathcal{A} to manipulate the clock in both directions, backward, and forward; in [3], only forward manipulation is allowed. Nonetheless, this does not result in a more powerful adversary.

access to real-world resources, functionality $\mathcal{F}_{\rm RO}$ as defined fails to limit the adversary on making a certain number of queries per round. Hence, we adopt a functionality wrapper [4,13] $\mathcal{W}(\mathcal{F}_{\rm RO})$ that wraps the corresponding resource to capture such restrictions. We highlight that our wrapper $\mathcal{W}(\mathcal{F}_{\rm RO})$ improves on previous wrappers in two aspects, in order to provide a more natural model of the real world: (1) We capture the pre-mining stage by letting the adversary query the RO with no restrictions (albeit polynomially bounded) before the CRS is released; (2) The wrapper limits adversarial access per nominal round by bounding the total number of queries that \mathcal{A} can make. The second aspect allows us to dispose the "flat" computational model and define the computational power in terms of the number of RO queries per round, which makes it possible to further refine the notion of a "respecting environment" (see below) that is suited for imperfect local clocks.

Diffusion. We adopt the peer-to-peer communication functionality $\mathcal{F}_{\text{Diffuse}}^{\Delta}$ (cf. [3]), which guarantees that an honestly sent message will be delivered to all the protocol participants within Δ rounds. Moreover, for those adversarially generated messages, $\mathcal{F}_{\text{Diffuse}}^{\Delta}$ forces them to be delivered to all the honest parties within Δ rounds after they are learnt by at least one honest participant. This captures the natural behavior of honest parties that they will forward all the messages that they have not yet seen to their peers.

We refer to [14] for a detailed description of the above functionalities.

2.3 Dynamic Participation

The notion of a "respecting environment" was introduced in [11] to model the varying number of participants in a protocol execution. In [2,3], the notion of *dynamic participation* was introduced aiming at describing the protocol execution in a more realistic fashion. Here we present a further refined classification of possible *types* of honest parties. See Table 1.

	Basic types of <i>honest</i> parties	
Resource	Resource unavailable	Resource available
random oracle $\mathcal{F}_{\mathrm{RO}}$	stalled	operational
network $\mathcal{F}_{\text{Diffuse}}^{\Delta}$	offline	online
${ m clock}\; \mathcal{F}_{ m ILCLOCK}$	time-unaware	time- $aware$
synchronized state	desynchronized	synchronized

Table 1. A classification of protocol participants.

Consider an honest party P at a given point of the protocol execution. We say P is *operational* if P is registered with the random oracle \mathcal{F}_{RO} ; otherwise, we say it is *stalled*. We say P is *online* if P is registered with the network; *offline* otherwise. We say P is *time-aware* if P is registered with the *imperfect* clock functionality $\mathcal{F}_{ILCLOCK}$, and *time-unaware* otherwise.

Further, we say P is synchronized if P has been participanting in the protocol for sufficiently long time and achieves a "synchronized state" as well as a "synchronized clock." "Synchronized clock" means P holds a chain that shares a common prefix (cf. [10]) with other synchronized parties; "synchronized clock" refers to that P maintains a local clock with time close to other synchronized parties. Otherwise, P is desynchronized. Additionally, P is aware of whether it is synchronized or not, and maintains a local variable isSync serving as an indicator for other actions.

Based on the above classification, we now define the notion of *alert* parties:

$alert \stackrel{\text{def}}{=} operational \land online \land time-aware \land synchronized.$

In short, alert parties are those who have access to all the resources and are synchronized; this requires them to join the protocol execution passively for some period of time. They constitute the core set of parties that carry out the protocol.

In addition, we define *active* parties to include all parties that are alert, adversarial, and time-unaware.

$active \stackrel{\text{def}}{=} alert \lor adversarial \lor time-unaware.$

Respecting environment in terms of computational power. Next, we provide the following generalization of "respecting environment" to relate it to computational power as opposed to number of parties. Our assumption is that during the whole protocol execution, the honest computational power is higher than the adversarial one (cf. the "honest majority" condition in [10] and follow-ups). The computational power is captured by counting the number of RO (hash) queries that parties make in each round. Further, we restrict the environment to fluctuate the number of such queries in a certain limited fashion.

Definition 2. For $\gamma \in \mathbb{R}^+$ we call the sequence $(h_r)_{r \in [0,B)}$, where $B \in \mathbb{N}$, (γ, s) -respecting if for any set $S \subseteq [0, B)$ of at most s consecutive integers, $\max_{r \in S} h_r \leq \gamma \cdot \min_{r \in S} h_r$.

We say that environment \mathcal{Z} is (γ, s) -respecting if for all \mathcal{A} and coins for \mathcal{Z} and \mathcal{A} the sequence of honest hash queries (h_r) is (γ, s) -respecting.

Note that the notion of respecting environment here is different from the "flat" model adopted in [10,11,12,4]. In a flat model, honest parties are assumed to have the same computational power, hence the total number of RO queries is a direct 1-to-1 map from the number of parties. The new respecting environment allows some subset of the honest parties to query the RO multiple times or stay stalled during a nominal round and hence it adapts to the "imperfect local clock" model used in this paper.

3 The Clock Synchronization Protocol with Public Setup

In this section we present the general approach and the various core building blocks of the new clock synchronization protocol—Timekeeper. For a complete description, refer to [14]. At a high level, Timekeeper is a Nakamoto-style PoWbased blockchain protocol together with time synchronization functionalities. Readers can think of it as a Bitcoin protocol with the following modifications:

- It replaces Bitcoin's original clock maintenance solution⁸ with a new clock synchronization scheme, which requires parties to use 2-for-1 PoWs [10] to mine and emit clock synchronization beacons and include them in an upcoming block. Furthermore, protocol participants will periodically adjust their local clock values based on the beacons collected in the blockchain and their (local) receiving time.
- Events are triggered by counting the number of local rounds (which is different from the convention that events in PoW-based blockchains are triggered by the arrival of blocks). In other words, the protocol has a clock synchronization *interval* of length R and a target recalculation *epoch* of length M that are defined in terms of the number of rounds; in addition, M is a multiple of R. Both of these values are hardcoded in the protocol. More precisely, parties will call the synchronization procedure (see Section 3.3) when their local clock enters round $\langle itvl, itvl \cdot R \rangle$ (this represents the last round in interval itvl; see below for details on the round structure); and for target in the next epoch they will call the target recalculation function (details see Section 3.4) when their local clock enters $\langle itvl, itvl \cdot R \rangle$, where (itvl mod (M/R)) = 0 (i.e., at the boundary of every (M/R) synchronization intervals).

See Figure 1 for an illustration of the protocol execution.



Fig. 1. An illustration of the clock synchronization protocol execution with one target recalculation epoch consisting of four clock synchronization intervals.

⁸ In Bitcoin's original implementation, miners will adjust their time based on three different sources: (1) their local system clock; (2) the median of clock values from peers; (3) the human operator (if the first two disagrees).

Next, we present the basic components that are employed in Timekeeper.

3.1 **Timekeeper** Timestamps

As opposed to the conventional approach where blocks' timestamps are integer values, timestamps (both blocks' and beacon values) in Timekeeper are represented by a pair of values interval number and round number $\langle \texttt{itvl}, \texttt{r} \rangle \in \langle \mathbb{N}^+, \mathbb{N}^+ \rangle$. Note that (ideally) one synchronization interval would last for R rounds (i.e., rounds $((i-1) \cdot \mathbb{R}, i \cdot \mathbb{R}]$ would belong to the *i*-th interval). However, in Timekeeper we let the lower bound be 0, which means that timestamps with a somewhat small round number are still valid. Specifically, a timestamp $(\texttt{itvl}, \texttt{r}) \triangleq \texttt{r} \leq \texttt{itvl} \cdot \mathbb{R}$. We note that this new treatment allows for some small distortion at the end of each interval—i.e., the round number of a few blocks at the beginning of the next interval may be smaller than the last block of the previous interval (we call these "retorted" timestamps); see Figure 2.



Fig. 2. An illustration of a segment of the blockchain with synchronization interval length R = 100. Blocks can have timestamp values equal to blocks in the previous interval.

Consider a chain of blocks in Timekeeper. Their timestamps should increase monotonically in terms of their interval number, and the round number in a single interval should also increase monotonically. More specifically, given two timestamps $\langle itvl_i, r_i \rangle$, $\langle itvl_j, r_j \rangle$ of two blocks $\mathcal{B}_i, \mathcal{B}_j$ respectively, if \mathcal{B}_i is an ancestor block of \mathcal{B}_j , they should satisfy the following predicate:

$$\begin{split} \mathsf{validTimestampOrder}(\langle \mathtt{itvl}_i, \mathtt{r}_i \rangle, \langle \mathtt{itvl}_j, \mathtt{r}_j \rangle) \\ & \triangleq \begin{cases} \mathsf{validTimestamp}(\mathtt{itvl}_i, \mathtt{r}_i) \land \mathsf{validTimestamp}(\mathtt{itvl}_i, \mathtt{r}_i) \\ & \land [(\mathtt{itvl}_i \leq \mathtt{itvl}_j) \lor (\mathtt{itvl}_i = \mathtt{itvl}_j \land \mathtt{r}_i < \mathtt{r}_j)] \end{cases} \end{cases} \end{split}$$

Furthermore, we will overload the notation of comparison operators based on the valid order of timestamps. E.g., "=" will denote that two timestamps are identical, and $\langle itvl_1, r_1 \rangle < \langle itvl_2, r_2 \rangle$ if and only if validTimestampOrder ($\langle itvl_1, r_1 \rangle, \langle itvl_2, r_2 \rangle$) holds. Other operators $>, \leq, \geq, \neq$ are defined similarly.

We also redefine "+, -" to describe the timestamp that is $k \in \mathbb{N}$ rounds before (resp., after) $\langle \texttt{itvl}, \texttt{r} \rangle$. Regarding addition, $\langle \texttt{itvl}, \texttt{r} \rangle + k = \langle \max\{\texttt{itvl}, \lceil(\texttt{r}+k)/\texttt{R}\rceil\}, \texttt{r}+k \rangle$. Intuitively, the additive operation simply adds k to r, and only increments itvl when it is going to become invalid. For subtraction, $\langle \texttt{itvl}, \texttt{r} \rangle - k = \langle \max\{1, \lceil(\texttt{r}-k)/\texttt{R}\rceil\}, \max\{1, \texttt{r}-k\}\rangle$. In other words, regarding the subtraction

operation, we only apply the operation on round, and the interval number is derived from the round after calculation. It does not output timestamps that are not "normally" belong to an interval. In case we do the subtraction operation for $k \ge \mathbf{r}$, it will return $\langle 1, 1 \rangle$.

Timekeeper's new approach to timestamps raises questions regarding the "trimming" of blockchains by counting the number of rounds. Recall that in [10] the notation $\mathcal{C}^{\lceil k}$ represents the chain that results from removing the k rightmost blocks. In this paper, we overload this notation to denote the chain that results from removing blocks with timestamps in the last k rounds with respect to the current time. Specifically, for $\mathcal{C} = \mathcal{B}_1 \mathcal{B}_2 \dots \mathcal{B}_n$ and local time $\langle itvl, r \rangle$, $\mathcal{C}^{\lceil k} = \mathcal{B}_1 \mathcal{B}_2 \dots \mathcal{B}_m$ is the longest chain such that $\forall \mathcal{B} \in \mathcal{C}^{\lceil k}$, Timestamp($\mathcal{B}) < \langle itvl, r \rangle - k$. In other words, \mathcal{B}_{m+1} is the first block (if it exists) such that Timestamp($\mathcal{B}_{m+1} \geq \langle itvl, r \rangle - k$ holds.

3.2 2-for-1 Proofs of Work and Synchronization Beacons

2-for-1 PoW is a technique that allows protocols to utilize a single random oracle $H(\cdot)$ to compose two separate PoW sub-procedures involving two distinct and independent random oracles $H_0(\cdot), H_1(\cdot)$. It was first proposed in [10] in order to achieve a better/optimal corruption threshold (from one-third to one-half) for the solution of the traditional consensus problem using a blockchain.

We refer to [10] for more details, and here we present a simple implementation with the clock synchronization application in mind. In order to do the 2-for-1 mining, a party P prepares a composite input w that is a concatenation of two inputs w_0, w_1 of two different sub-procedures S_0, S_1 , respectively. I.e., $w = w_0 || w_1$. After selecting a nonce ctr, quering the random oracle with ctr || w and getting result u, P checks if u < T which implies success in sub-procedure S_0 ; P also checks if $[u]^{\mathbb{R}} < T$ (where $[u]^{\mathbb{R}}$ denotes the reverse of a bitstring u) which indicates success in sub-procedure S_1 . After successfully generating a PoW for S_0 (resp., S_1), in order to let parties others check validity, the proof will include the nonce and the entire composite input ctr || w. Note that sub-procedure S_0 (resp., S_1) only cares about its corresponding part w_0 (resp., w_1), and treat the other part as dummy information.

The 2-for-1 PoW technique has several advantages when compared with the straightforward approach that would simply utilize two different random oracles. The most prominent advantage is that it prevents the adversary \mathcal{A} from concentrating its computational power on one RO and thus gain advantage in the corresponding sub-procedure.

Synchronization beacons. In addition to the conventional blocks constituting the blockchain, protocol participants in Timekeeper also produce another type of "tiny" blocks using 2-for-1 PoWs. We call these blocks *clock synchronization beacons* ("beacons" for short) since they are used to report parties' local time and synchronize their clocks.

In more detail, one clock synchronization beacon SB is a tuple with the following structure.

$$SB \triangleq \langle \langle itvl, r \rangle, P, \eta_{itvl}, ctr, blockLabel \rangle,$$

where $\langle itvl, r \rangle$ is the local time SB reports; P denotes the identity of its miner; η_{itvl} is some fresh randomness in the current interval; ctr represents the nonce of the PoW and *blockLabel* is the associated block input. Note that SB must record the identity of its miner because there might be multiple beacons, mined by different parties, reporting the same timestamp as well as nonce value; otherwise, it would be impossible for the parties to distinguish such beacons. Worse still, other participants would not be able to distinguish the same beacon SB when they receive SB multiple times. Regarding η_{itvl} , it is a string associated with interval itvl for the purpose of preventing the adversary \mathcal{A} from mining beacons with future timestamps. In other words, protocol participants (including \mathcal{A}) can only learn η_{itvl} after they have (almost) finished interval itvl - 1. We present the structure of intervals in detail and how we compute η_{itvl} in Section 3.3 and treat it as a communal bitstring here. We note that parties can learn η_{itvl} from their local chain, and indeed SB does not need to include η_{itvl} (P can prune those beacons that are invalid with η_{itvl} in their local view). We keep η_{itvl} in the description for convenience.

Regarding the structure of a blockchain block \mathcal{B} , we adopt the similar structure as inin [11] (with the dummy information in the 2-for-1 PoWs):

$$\mathcal{B} \triangleq \langle h, \mathtt{st}, \langle \mathtt{itvl}, \mathtt{r} \rangle, ctr, txLabel \rangle,$$

where h is the reference to the previous block, st the Merkle root of the block content, $\langle itvl, r \rangle$ its timestamp, *ctr* the nonce of PoW, and *txLabel* the binded beacon input.

We are now ready to describe how the parties in Timekeeper do the 2-for-1 PoW mining. The composite input prepared in Timekeeper is different from the trivial instance above, in that the term $\langle \texttt{itvl}, \texttt{r} \rangle$ appears in both blocks and beacons. Hence, simply concatenating two inputs introduces redundant information in the PoW. When a party P is ready to perform the mining procedure, P binds the nonce, the blocks' input and beacon input together as

$$\langle ctr, h, \mathtt{st}, \langle \mathtt{itvl}, \mathtt{r} \rangle, \mathsf{P}, \eta_{\mathtt{ep}}^{\mathcal{C}} \rangle$$

and hand them over to random oracle \mathcal{F}_{RO} . Let u denote the result from \mathcal{F}_{RO} . If u < T (i.e., the block query succeeds), P finds a new block $\mathcal{B} = \langle h, \mathsf{st}, \langle \mathsf{itvl}, \mathsf{r} \rangle$, $ctr, txLabel \rangle$ where $txLabel := \langle \mathsf{P}, \eta_{\mathsf{ep}}^{\mathsf{C}} \rangle$; if $[u]^{\mathsf{R}} < T$ (the beacon query succeeds), P gets a new beacon SB = $\langle \langle \mathsf{itvl}, \mathsf{r} \rangle, \mathsf{P}, ctr, blockLabel \rangle$, where $blockLabel := \langle h, \mathsf{st} \rangle$. Note that for the sake of presentation, we reorder the content of blocks and beacons so that they are inconsistent with the input to the PoW.

After receiving the result from \mathcal{F}_{RO} , P checks if it was able to successfully generate a new block. In addition, P checks if he successfully produces a beacon

but only when P's local clock stays in the *beacon mining and inclusion* phase. Namely, P reports a timestamp that satisfies a certain criterion (details in Section 3.3).

3.3 Clock Synchronization Intervals and the Synchronization Procedure

As mentioned earlier, Timekeeper participants will periodically adjust their local clock. We call the time interval between two adjustment points⁹ a *clock synchronization interval* (or "interval" for short). Ideally, one interval will last for R rounds. The actual number of local rounds that parties observe may differ according to the shift computed in the previous interval (we will show later that the shift computed in every interval is well-bounded). When party P's local clock gets to the last round of an interval, it will call the synchronization procedure (see below), which adjusts its local clock and gets the fresh randomness to run the next interval.

Interval Structure. Timekeeper divides one interval into three phases: (1) view convergence, (2) beacon mining and inclusion and (3) beacon-set convergence. The phase parties stay in depends on their local clocks. Furthermore, parties will keep track of the (local) arriving time of a synchronization beacon as long as it is online. In this section we describe these three phases as well as the bookkeeping function and explain the design intention behind them.

View convergence. When a party P's local clock reports a time $\langle itvl, r \rangle$ such that $r < (itvl-1) \cdot R + K$, P is in the view convergence phase. Note that this also includes rounds with potentially retorted timestamps. In this phase, if P is alert, it will try to mine the next block with the 2-for-1 PoW technique (i.e., the input information that P forwards to the \mathcal{F}_{RO} functionality does not need to be changed); nonetheless, P will not check if he successfully mines a beacon after P acquires the output. This is because all the beacons obtained in this phase are invalid in that they report an undesirable timestamp.

The general motivation for introducing the view convergence phase and letting parties wait for some period of time at the beginning of an interval is that we would like parties to start mining beacons with a *consistent* view of the previous interval. Since K is larger than the common prefix parameter (we will quantify K in later, in Section 4.2), at the end of the view convergence phase of interval itvl+1, alert parties will have a common view of interval itvl. In other words, they will agree on all the blocks in interval itvl, and the adversary \mathcal{A} will not be able to apply any changes to these blocks. Hence, alert parties agree on the number of blocks in the previous interval, which decides the mining difficulty within the current interval. (This will used in our new target recalculation function, presented in Secion 3.4.) Parties will mine beacons with the same difficulty,

⁹ The first interval in particular lies between the beginning of the execution and the first time parties adjust their clock.

and this simplifies the protocol description as well as its analysis. Furthermore, alert parties will compute the same fresh randomness as

$$\eta_{\mathtt{itvl}+1} \triangleq G(\eta_{\mathtt{itvl}} \parallel (\mathtt{itvl}+1) \parallel v), \tag{1}$$

where v is the concatenation of all block hashes in interval *itvl*. Note that we adopt a different hash function $G(\cdot)$ (as opposed to $H(\cdot)$) to compute the next fresh randomness that is not used in the 2-for-1 PoW, which does not consume any queries to random oracle \mathcal{F}_{RO} .

Recall that by assumption the adversary \mathcal{A} has full knowledge of the network, and hence it can learn all honest blocks from the previous interval immediately and manipulate the chain at will for up to a number of rounds bounded by the common prefix parameter, allowing \mathcal{A} to mine the synchronization beacons before the alert parties start to mine. We call this period where \mathcal{A} starts ahead of time the *pre-mining* stage. Nonetheless, we will show later that there will be at least one block generated by an alert party near the end of interval *itvl*, which prevents the adversary from pre-mining for too long a time.

Beacon mining and inclusion. When a party P's local clock is in rounds $\langle itvl, r \rangle$ satisfying $(itvl - 1) \cdot R + K \leq r \leq itvl \cdot R - K$, P is in the beacon mining and inclusion phase. Next, we define the predicate $I_{sync}(itvl)$ to extract the set of timestamps in this phase. Formally,

$$I_{\text{sync}}(\text{itvl}) \triangleq \{(\text{itvl}-1) \cdot \mathbf{R} + \mathbf{K}, \dots, \text{itvl} \cdot \mathbf{R} - \mathbf{K}\}.$$
 (2)

For convenience, we slightly overload this predicate. When the input is a timestamp, $I_{sync}(\langle itvl, r \rangle)$ outputs whether $\langle itvl, r \rangle$ stays in a beacon mining and inclusion phase. I.e., $I_{sync}(\langle itvl, r \rangle) = true$ if $r \in I_{sync}(itvl)$, and false otherwise.

After entering this phase, P will use a 2-for-1 PoW to mine both blocks and clock synchronization beacons. During interval itvl, the output will be a beacon which indicates its local time and value $SB \triangleq \langle \langle itvl, r \rangle, P, ctr, blockLabel \rangle$. Regarding the mining difficulty, Timekeeper will set the same target value for blocks and beacons¹⁰. In other words, the expected number of blocks and of beacons in this phase are equal.

After a beacon is successfully generated, it will be diffused into the network via $\mathcal{F}_{\text{Diffuse}}^{\text{sync}}$. P will include a beacon SB into the pending block content if SB is valid w.r.t. the current interval. Next, we describe how they check the validity of a beacon is checked. The format of a beacon SB with respect to interval itvl is correct if and only if it reports a timestamp $\langle \text{itvl}, \mathbf{r} \rangle$ such that $\mathbf{r} \in I_{\text{sync}}(\text{itvl})$. We say a beacon SB is *valid w.r.t. chain* C if and only if its format is correct and the hash value of SB (after concatenating with the fresh randomness in C) is smaller than the corresponding mining target. P will try to include all the

¹⁰ We will adopt the same target for simplicity. Indeed, maintaining a constant ratio between the difficulty level of blocks and that of beacons will work.

(valid) beacons mined in the current interval itvl with timestamps earlier than the current local time but which have not yet been included in the blockchain. Specifically, at round $\langle itvl, r \rangle$, all valid beacons recording timestamp $\langle itvl, u \rangle$ with $u \leq r$ will get into P's pending block content.

When P's local clock goes past the last round of beacon mining and inclusion phase, it stops checking the beacon hash output and it no longer includes beacons in the next block. Beacons that are generated and diffused right at the end of this phase get dropped.

Beacon-set convergence. The third and last phase—beacon-set convergence consists of the last K rounds in an interval. In other words, a party P is in this phase when P reports a timestamp $\langle itvl, r \rangle$ with $r > itvl \cdot R - K$. During this phase, P behaves similar to the first phase. I.e., it will not check for the 2-for-1 PoW result to see if the beacon generation succeeds.

Parties have to wait for at least K rounds to ensure that they share a *consistent* view of the set of beacons included in the current interval (except with some negligible probability). This phase cannot be omitted since only when parties agree on the same beacon set can the synchronization procedure maintain the protocol's security properties (Section 4).

Beacon arrival booking. In order to adjust its clock, P also needs the local receiving time of all beacons that have been included in the chain. Hence, P will maintain a local registry that records the beacons it receives as well as their arrival time. More specifically, this local beacon ledger is an array of synchronization beacons. For each beacon SB, a pair $(a, flag) \in \langle \mathbb{N}^+, \mathbb{N}^+ \rangle \times \{ \text{final}, \text{temp} \}$ is assigned to it. Consider a round $\langle \text{itvl}, \mathbf{r} \rangle$ when P receives a beacon SB with Timestamp(SB) = $\langle \text{itvl}', \mathbf{r}' \rangle$.

- If itvl' ≤ itvl, which means the beacon SB is generated in the current or previous interval¹¹. P will drop SB if it is not valid w.r.t. its localchain; otherwise, P will assign ((itvl, r), final) to SB. This means that all the information gathering regarding this beacon has been finalized and it is ready to be used.
- If itvl' > itvl, the beacon is generated in the future. P will assign ((itvl,r), temp) to SB, which indicates that modifications on the receiving time may be applied in the future. Note that parties may not know the fresh randomness in future intervals (for example, if they are newly joint parties and have not yet synchronized with the blockchain or they are alert but receive forthcoming beacons). Hence they cannot check the validity of beacons with temp flag. Nevertheless, invalid beacons would be excluded from the registry after P learns the upcoming fresh randomness.

¹¹ Beacons generated in previous intervals are stale in that P has already passed the synchronization point associated with these beacons, and they will never be used in the future. We list them for completeness.

If P receives multiple beacon messages with the same creator and time reported, P will adopt the first one it receives as its arrival time.

The Synchronization Procedure. At the end of an interval (i.e., when the local time reports $r = itvl \cdot R$), parties will use the beacons information to compute a value shift that indicates how much the logical clock should be adjusted. (See [14] for the complete specification.)

Adjusting the local clock. When a party P's local clock reaches round $\langle itvl, itvl, R \rangle$ and P has finished the round's regular mining procedure, P will adjust its local clock based on the beacons recorded on chain and their local receiving time. More specifically, P will extract all the beacons from the beacon mining and inclusion phase, and compute the differences between their timestamp and local receiving time Timestamp(SB) – arrivalTime(SB). Since the timestamp of SB and its arrival time share the same interval index, we only need to compute the difference between their round numbers. Subsequently, all the beacons will be ordered based on this difference and a shift will be computed by selecting the median difference therein. Formally,

$$\mathsf{shift}_{\mathsf{itvl}}^{\mathsf{P}} \triangleq \mathsf{med}\{\mathsf{Timestamp}(\mathsf{SB}) - \mathsf{arrivalTime}(\mathsf{SB}) \mid \mathsf{SB} \in \mathcal{S}_{\mathsf{itvl}}^{\mathsf{P}}\}.$$
 (3)

In case there are two median beacons SB_1, SB_2 , parties will adjust $shift_{itv1}^{P} \triangleq [(Timestamp(SB_1) - arrivalTime(SB_1) + Timestamp(SB_2) - arrivalTime(SB_2))/2].$ Afterwards, P will update its local clock to $\langle itvl + 1, r + shift \rangle$. Later we show that this update strategy in the synchronization procedure allows parties' clocks to remain in a narrow interval and do not deviate too much from the nominal time.

Note that parties will enter local round (itvl, r) where $r = itvl \cdot R$ only once. If they enter some time (itvl', r) in the future, we will get itvl' > itvl and they will never revert back.

Mining with backward-set clocks. After the adjustment at the end of intervals, and shift is added to P's local clock, it may set its local time to values $\langle itvl, r \rangle$ such that $r \leq (itvl-1) \cdot R$ (i.e., the retortion effect that was mentioned earlier). Nonetheless, P can continue to mine blocks with this timestamp and its local clock will eventually proceed to a time value of regular format (i.e., $r > (itvl-1) \cdot R$).

We compare this treatment with the similar scenario in a PoS blockchain [3]. In [3], setting local clocks backward is never a problem since parties can keep silent during this period. Due to the nature of PoS-based blockchains, parties do not need to do anything if they are not assigned the leader slot. In our context, however, adopting the same 'silence' policy contradicts the basic nature of PoWbased blockchains as parties will forfeit the chance to extend their local chain. In other words, there is no point for an activate party to not make RO queries. This is taken care of by Timekeeper's timestamping scheme.

Updating the beacon arrival time registry. Notice that the beacon information stored in a party P's arrival time registry is closely related to which interval P stays in; after P enters the next interval, it needs to update the beacon bookkeeping. P will apply a shift computation for all beacons with flag temp. Furthermore, for those beacons that report a timestamp with interval equal to the incoming one, their flag will be set to final. In more detail, at the end of interval itvl, for all eligible SB in the beacon registry, their associated pair ($\langle itvl_{SB}, r_{SB} \rangle$, temp) will be updated to ($\langle itvl_{SB}, r_{SB} + shift \rangle$, final) if $itvl_{SB} = itvl + 1$. Note that for those beacons whose flags are set to final, P will removed all invalid ones from the registry after the update.

3.4 The Target Recalculation Function

If the mining target is not set appropriately ("appropriately" means that the block generation rate according to the current hashing power and target is somewhat steady; see [11]), PoW-based blockchain protocols fail to maintain any of the security properties in a permissionless environment. In Bitcoin, the target is adjusted after receiving the last block of the current epoch (and an epoch consists of 2016 blocks). Based on the time elapsed to mine these blocks, a new target is set based on the previous target value and the variation is proportional to the time elapsed. Note that Bitcoin's target recalculation function is not the only way to adjust the difficulty level. A large number of other recalculation functions have been proposed in alternate blockchains (e.g., Ethereum, Bitcoin Cash, Litecoin), with their security asserted by either theoretical analysis or empirical data.

In Timekeeper, we propose a new target recalculation function that is suitable for the new setting. Intuitively, our function is a reversed version of Bitcoin's original function, namely, protocol participants wait for some fixed number of rounds M (in their local view) to update the difficulty level. We call such M number of rounds a *target recalculation epoch*. Moreover, Timekeeper sets M as a multiple of R, which makes the target recalculation epoch consist of several clock synchronization intervals, and the start and end point of an epoch coincide with the start and end of different synchronization intervals. Recall that in the Timekeeper timestamp scheme introduced in Section 3.1, the first term in $\langle itvl, r \rangle$ does not directly reflect which target recalculation epoch it is in. For simplicity, we introduce function TargetRecalcEpoch that maps the protocol timestamp to the target recalculation epoch it belongs to:

 $\mathsf{TargetRecalcEpoch}(\langle \mathtt{itvl}, \mathtt{r} \rangle) \triangleq [\mathtt{itvl}/(M/R)].$

In addition, we introduce a function EpochBlocks which extracts all the blocks in chain C that belong to target recalculation epoch ep. Formally, given ep ≥ 1 ,

 $\mathsf{EpochBlocks}(\mathcal{C}, \mathsf{ep}) \triangleq \{\mathcal{B} : \mathcal{B} \in \mathcal{C} \land \mathsf{TargetRecalcEpoch}(\mathsf{Timestamp}(\mathcal{B})) = \mathsf{ep}\}.$

Also for convenience, we let $\mathsf{EpochBlockCount}$ be a function that returns the number of blocks in chain \mathcal{C} that belong to epoch ep. We also extend the input

domain of epoch numbers to 0 and let it output Λ_{epoch} (the ideal number of blocks) to capture the fact that the target at the beginning of an execution is set appropriately and hence maintains the ideal block generation rate. Formally,

$$\mathsf{EpochBlockCount}(\mathcal{C}, \mathsf{ep}) \triangleq \begin{cases} |\mathsf{EpochBlocks}(\mathcal{C}, \mathsf{ep})| & if \ \mathsf{ep} \ge 1\\ \Lambda_{\mathsf{epoch}} & if \ \mathsf{ep} = 0 \end{cases}$$
(4)

Going back to the algorithm, for the first epoch (ep = 1) parties will adopt the target value of the genesis block (T_0). I.e., $T_1 = T_0$. Regarding other epochs (ep > 1), parties will figure out how many blocks are produced in the previous epoch, and set the next target based on the previous one. This variation is proportional to the ratio of expected number of blocks Λ_{epoch} and the actual number. I.e., for epoch ep + 1,

$$T_{\text{ep}+1} \triangleq \frac{\Lambda_{\text{epoch}}}{\Lambda} \cdot T_{\text{ep}}, \quad \text{ep} \in \mathbb{N}^+,$$
(5)

where Λ is the number of blocks in epoch ep—in other words, the size of EpochBlocks(C, ep).

In order to prevent the "raising difficulty attack" [6], the maximal target variation in a single recalculation step still needs to be bounded (we denote this bound by τ). Specifically, if $\Lambda > \tau \cdot \Lambda_{\text{epoch}}$, $T_{\text{ep}+1}$ will be set as T_{ep}/τ ; on the other hand, if $\Lambda < \Lambda_{\text{epoch}}/\tau$, $T_{\text{ep}+1}$ will be set as $\tau \cdot T_{\text{ep}}$.

Remark 1. We observe that, compared to the Bitcoin case, the adversary \mathcal{A} in Timekeeper is in a much worse position to carry out the raising difficulty attack. This is because in Bitcoin, in order to significantly raise the difficulty in the next epoch, \mathcal{A} only needs to mine 2016 blocks with close timestamps; in the case of Timekeeper, however, the adversary has to mine $\tau \cdot \Lambda_{epoch}$ blocks (with fake timestamps) in order to raise the same level of difficulty. The number of blocks that \mathcal{A} needs to prepare is τ times larger than that in Bitcoin (assuming both protocols share the same number of expected blocks in an epoch).

3.5 Newly Joining Parties

Recall that Timekeeper runs in a permissionless environment where parties can join and leave at will. As such, it is essential that newly joining parties can learn the protocol time to become alert and participate in the core mining process. More specifically, after the joining procedure, newly joining party P's local clock should report a time in a sufficiently narrow interval with all other alert parties, at which point P can claim also being alert.

Based on the fine-grained classification of types of parties in our dyanmic participation model (Section 2.3), newly joining parties can be classified into two types: (1) parties that are temporarily de-registered from $\mathcal{F}_{\rm RO}$, and (2) parties that start with bootstrapping from the genesis block, or parties that temporarily

lose the network connection (i.e., de-registered from $\mathcal{F}_{\text{Diffuse}}^{\Delta}$), or parties that are temporarily de-registered from $\mathcal{F}_{\text{ILCLOCK}}$.

For parties that are stalled for a while, since they do not miss any clock tick or other necessary information from the network, they can easily re-join by simulating clock adjustments locally (details see [14]). For the rest of newly joining parties, they will be classified as de-synchronized (note that parties are aware of their synchronization status), and will run the joining procedure JoinProc, which we now describe.

Procedure JoinProc. In order to synchronize its clock, a newly joining party P needs to "listen" to the protocol for sufficiently long time. We describe the joining process below, which is similar to that in [3]. The main difference is that we adopt the heaviest-chain selection rule in order to adapt to the PoW context. The complete specification of this protocol is presented in [14], and the default parameters values are summarized in Table 2.

Parameter	Default	Phase
$t_{\rm off}$	2K	В
$t_{\rm gather}$	5R/2	С
$t_{\rm pre}$	3K	D

Table 2. Parameters of the joining procedure and their corresponding phases.

- Phase A (state reset). When all resources are available to P, after resetting all its local variables, P invokes the main round procedure triggering the join procedure.
- Phase B (chain convergence, with parameter t_{off}). In the second activation upon a MAINTAIN-LEDGER command, the party will jump to phase B and stay in phase B for t_{off} rounds. During this phase, the party applies the *heaviest-chain selection rule* maxvalid to filter its incoming chains. The motivation behind Phase B is to let P build a chain that shares a sufficiently long common prefix with all alert parties. Note that since P has not yet learnt the protocol time, it cannot filter out chains that should be put aside in the futureChains. Hence, the chain held by P may still contain a long suffix built entirely by the adversary. However, it can be guaranteed that this adversarial fork can happen for up to k rounds ahead. Thus, the beacons recorded before the fork can be used to compute the adjustment and their local arrival times will be reliable.
- Phase C (beacon gathering, with parameter t_{gather}). Once a party P has finished Phase B, it continues with Phase C, the beacon-gathering phase. During this phase, P continues to collect and filter chains as in Phase B. In addition, P now processes and bookkeeps the beacons received from $\mathcal{F}_{\text{Diffuse}}^{\text{sync}}$. At a high level, this phases' length parameter t_{gather} guarantees that: (1) enough beacons are recorded to compute a reliable time shift; (2)

enough time has elapsed so that the blockchain reaches agreement on the set of (valid) beacons to use. At the end of Phase C, P is able to reliably judge valid arrival times.

Phase D (shift computation, with parameter t_{pre}). Since party P has now built a blockchain sharing a common prefix with any alert party, and has bookkeeped synchronization beacons for a sufficiently long time, P starts from the earliest interval i^* such that (1) the arrival times of all beacons included in blocks within the beacon mining and inclusion phase of interval i^* have been locally bookkeeped, and (2) all of these beacons arrived sufficiently later than the start of Phase C (parameterized by $t_{\rm pre}$ rounds). Based on this information, P computes the shift value as alert parties do at the boundary of synchronization interval i^* . P concludes Phase D when the adjusted time is a valid timestamp in interval $i^* + 1$ (in other words, r does not exceed $(i^* + 1)R$; otherwise, P updates the local arrival time of beacons with flag temp and repeats the above process with interval $i^* + 1$. We note that if Phase D involves the computation w.r.t. multiple intervals, the local time may temporarily be set as an invalid timestamp. Nevertheless, eventually after P has passed (2K+5R/2) (local) rounds, P will end up with a valid timestamp that with overwhelming probability is close enough to those of all alert parties.

4 Protocol Analysis

Our ultimate goal is to show that, at any point of the protocol's execution, the timestamps reported by all alert protocol participants of Timekeeper will satisfy the properties defined in Definition 1. We start off with some additional definitions and preliminary results.

4.1 Notation, Definitions and Preliminary Propositions

We note that several of the analytical tools proposed in [11,12] do not directly apply in the environment (with $\mathcal{F}_{ILCLOCK}$) where Timekeeper runs in. Therefore, we first extend and enhance these tools to adapt them to this new environment.

Our probability space is over all executions of length at most some polynomial in κ and λ ; we use **Pr** to denote 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.

For the sake of convenience, we define a *nominal time* that coincides with the internal variable time in $\mathcal{F}_{\text{ILCLOCK}}$. Recall that time aims at recording how many times the functionality sends clock ticks to all registered honest parties.

Definition 3 (Nominal Time). Given an execution of Timekeeper, any prefix of the execution can be mapped deterministically to an integer r, which we call

nominal time, as follows: r is the value of variable time in the clock functionality at the final step of the execution prefix which is obtained by parsing the prefix from the genesis block and keeping track of the honest party set registered with the clock functionality (bootstrapped with the set of inaugural alert parties). (In case no honest party exists in the execution, r is undefined).

Note that we adopt r to denote the nominal time, which is different from the protocol timestamp $\langle itvl, r \rangle$.

If at a nominal round r exactly h parties query the oracle with target T, the probability of at least one of them will succeed is

$$f(T,h) = 1 - (1 - pT)^h \le pTh$$
, where $p = 1/2^{\kappa}$.

During nominal round r, alert 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. Moreover, the initial target T_0 implies in our model an initial estimate of the number of honest RO queries h_0 ; specifically, $h_0 = 2^{\kappa} \Lambda_{\text{epoch}}/(T_0 M)$, i.e., the number of parties it takes to produce Λ_{epoch} blocks of difficulty $1/T_0$ in time M. For convenience, we denote $f_0 = f(T_0, h_0)$ and simply refer to it as f. Also note that the ideal number of blocks $\Lambda_{\text{epoch}} = Mf$, so in the analysis we will use Mf to represent Λ_{epoch} .

"Good" properties. Next, we present some definitions which will allow us to introduce a few ("good") properties, serving as an intermediate step towards proving the desired clock properties.

Let us consider the boundary of two target recalculation epochs. Recall that Bitcoin's target recalculation algorithm defines epoch in terms of the number of blocks (m blocks forms an epoch). Thus, a block with block height a multiple of m is the last block of an epoch. While it might be manipulated, its timestamp naturally becomes the proof that miners have adjusted their difficulty and entered the next epoch (known as the *target recalculation point* [11]). In contrast, **Timekeeper** adopts a new target recalculation function (see Section 3.4) that divides the epoch based on the parties' local view. While we can still define a target recalculation point based on one party's local view, parties can *never* agree on a point where they enter the next epoch based on nominal time.

In order to circumvent the above obstacle, we extend the notion of target recalculation point to target recalculation *zone*. See Figure 3 for an illustration. Intuitively, a "target recalculation zone" w.r.t. epoch **ep** is a sequence of consecutive nominal rounds such that during these nominal rounds, at least one alert party crosses its own target recalculation point w.r.t. epoch **ep**. For convenience, we assume a "safe" start—i.e., the first epoch also has a target recalculation zone, and it naturally satisfies all good properties we will later define.

Definition 4.

- Nominal time r is good if $f/2\gamma^2 \leq ph_r T_r^{\min}$ and $ph_r T_r^{\max} \leq (1+\delta)\gamma^2 f$.



Fig. 3. An illustration of the target recalculation zone $Z_{ep} = \{t, \ldots, t + 6\}$.

- Round (itvl, r) is a target-recalculation point w.r.t. epoch ep if (r = itvl \cdot R) \wedge [itvl mod (M/R) = 0].
- A sequence of consecutive nominal rounds $Z_{ep} = \{r\}$ is a target recalculation zone w.r.t. target recalculation epoch ep if during Z_{ep} some subset of synchronized parties are in the logical round that is a target recalculation point w.r.t. ep 1.
- A target-recalculation zone Z_{ep} is good if for all $h_r, r \in Z_{ep}$ the target T_{ep} satisfies $f/2\gamma \leq ph_r T_{ep} \leq (1+\delta)\gamma f$.
- A chain is good if all its target-recalculation zones are good.
- A chain is stale if for some nominal time u it does not contain an honest block computed after nominal time $u \ell 2\Delta 2\Phi$.
- The blocklength of an epoch ep on a chain C is the number of blocks in C with timestamp $(itvl, \cdot)$ such that TargetRecalcEpoch((itvl, r)) = ep.

We would like to prove that, at a certain nominal round r of the protocol execution, alert parties enjoy good properties on their local chains and reported timestamps. Towards this goal, we extract all chains that either belong to alert parties at r or have accumulated sufficient difficulty and thus might be adopted in the future. We denote this chain set by S_r :

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

Next, we define a series of useful predicates with respect to the potential chain set S_r and parties' local clocks at nominal round r. Note that Φ is a constant that is the ideal maximal skew of all alert clocks, and $\Phi = \Delta + \Phi_{clock}$ where Δ is the network delay and Φ_{clock} is the maximal clock drift that \mathcal{A} can set (see Section 2.1).

Definition 5. For a nominal round r, let:

- GOODCHAINS $(r) \triangleq$ "For all $u \leq r$, every chain in S_u is good."
- GOODROUND $(r) \triangleq$ "All rounds $u \leq r$ are good."
- NOSTALECHAINS $(r) \triangleq$ "For all $u \leq r$, there are no stale chains in S_u ."
- COMMONPREFIX $(r) \triangleq$ "For all $u \leq r$ and $C, C' \in S_r$, head $(C \cap C')$ was created after nominal round $u \ell 2\Delta 2\Phi$."
- BLOCKLENGTH(r) \triangleq "For all u < r and $C \in S_u$, the blocklength Λ of any epoch ep in C satisfies $\frac{1}{2(1+\delta)\gamma^2} \cdot mf \leq \Lambda \leq 2(1+\delta)\gamma^2 \cdot mf$
- GOODBEACONS(r) \triangleq "For all u < r and the beacon set S_{itvl} bookkeeped during any interval itvl, more than half of beacons within S_{itvl} are generated by honest parties".
- GOODSHIFT(r) \triangleq "For all u < r, and the alert party P_i that adjusts its local clock at round u, P_i computes shift_i that $-2\Phi \leq \mathsf{shift}_i \leq \Phi$ ".
- GOODSKEW $(r) \triangleq$ "For all alert parties in nominal time r, their local time in this round differs by at most Φ if they are in the same interval or differs by at most 2Φ if they are in different intervals." Formally,

$$GOODSKEW(r) :\Leftrightarrow \left(\forall \mathsf{P}_1, \mathsf{P}_2 \in \mathcal{P}_{\mathsf{alert}}[r] : \left| \begin{vmatrix} \mathsf{r}_1 - \mathsf{r}_2 \end{vmatrix} \le \Phi & \text{ if } \mathsf{itvl}_1 = \mathsf{itvl}_2 \\ |\mathsf{r}_1 - \mathsf{r}_2| \le 2\Phi & \text{ if } \mathsf{itvl}_1 \neq \mathsf{itvl}_2 \\ \end{vmatrix} \right)$$

where $\langle itvl_1, r_1 \rangle$ and $\langle itvl_2, r_2 \rangle$ are the timestamps that P_1 and P_2 reports during r^{12} .

Random variables and (Δ, Φ) -isolated success. Next, for the purpose of estimating the difficulty acquired by honest parties during a sequence of rounds, we define the following random variables w.r.t. nominal round r.

- D_r : the sum of the difficulties of all blocks computed by alert parties at nominal round r.
- Y_r : the maximum difficulty among all blocks computed by alert parties at nominal round r.
- Q_r : equal to Y_r when $D_u = 0$ for all $r < u < r + \Delta + \Phi$ and 0 otherwise.

We call a nominal round r such that $D_r > 0$ successful and one wherein $Q_r > 0$ isolated successful. An isolated successful round guarantees the irreversible progress of the honest parities.

We highlight that, under the imperfect local clock model $\mathcal{F}_{\text{ILCLOCK}}$, the notion of an "isolated successful round" needs to be re-considered as parties' local clocks may span some consecutive rounds. Assuming a Φ -drift is maintained during the sequence of rounds we are interested in, an irreversible contribution to the chain happens when the (nominal) distance between such success and the following success is at least $\Phi + \Delta$ rounds. This is because the block producer may have a local clock that is already Φ rounds behind other alert parties, and it takes Δ rounds to diffuse the block. This cancels out other parties' successes for up to

¹² If P passes multiple local rounds in nominal round r, we require that all of these timestamps should satisfy the predicate.

 $\Phi + \Delta$ rounds. As a result, we call such event a (Δ, Φ) -isolated successful, which is the for the new formulation of Q_r (cf. [12]). Note that this (Δ, Φ) -isolated successful round is meaningful only when the protocol is able to maintain a Φ -bounded skew during the sequence of rounds we are considering.

Recall that the total number of hash queries alert parties (resp., the adversary) can make during nominal round r is denoted by h_r (resp., t_r). For a sequence of rounds S we write $n(S) = \sum_{r \in S} n_r$ and similarly, t(S), D(S), Q(S).

Regarding the adversary \mathcal{A} , while \mathcal{A} may query the random oracle for an arbitrarily low target and obtain blocks with arbitrarily high difficulty, we wish to upper-bound the difficulty it can accrue during a set of J queries. Consider a set of consecutive adversarial queries J and associate it with the target of the first query (this target is denoted by T(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.

Let \mathcal{E}_{r-1} fix the execution just before (nominal) 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 h_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.

Blockchain properties. We use blockchain properties as formulated in [10,11] as an intermediate step towards proving the clock properties and achieve our blockchain synchronizer. Next, we briefly describe these properties: common prefix, chain growth, chain quality and existential chain quality.

Notably, we consider common prefix in terms of number of rounds. I.e., honest parties will agree on a settled part of the blockchain with timestamps at most a given number of rounds before their local time.¹³ Let $C^{\lceil k}$ denote the chain resulting from removing all rightmost blocks with timestamp larger than $\mathbf{r} - k$, where \mathbf{r} is the current (local) time. We can now define common prefix as follows.

- Common Prefix (with parameter $k \in \mathbb{N}$). For any two alert parties $\mathsf{P}_1, \mathsf{P}_2$ holding chains $\mathcal{C}_1, \mathcal{C}_2$ at rounds r_1, r_2 , with $r_1 \leq r_2$, it holds that $\mathcal{C}_1^{\lceil k} \preccurlyeq \mathcal{C}_2$.

Regarding chain growth, the lemma below provides a lower bound on the irreversible progress of achieved by the honest parties regardless of any adversarial behavior. This lemma has appeared in previous analyses under varying settings, evolving from the synchronous network and static environment ([10]), to a dynamic environment ([11]), and further to a bounded-delay network set-

¹³ While most of the previous work considers common prefix in terms of number of blocks, we note that these two definitions are equivalent. This is due to the fact that if the protocol guarantees security, then the block generation rate is somewhat steady (cf. [11]) and thus the number of blocks generated during a period of time can be inferred from its length and the highest mining speed.

ting ([12]). The next lemma extends the chain growth property to a Δ -bounded network delay, Φ -bounded clock drift and dynamic environment.

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

4.2 Protocol Parameters and Their Conditions

We summarize all Timekeeper parameters in Table 3 in Appendix A. It is worth noting that ϵ is a small constant regarding the quality of concentration of random variables (it will appear in the typical executions in Section 4.3). We introduce a parameter λ —which is related to the properties of the protocol—to simplify several expressions. Protocol parameter λ and the RO output length κ are the seucrity parameters of Timekeeper.

In order to get desired convergence and perform meaningful analysis, we consider a sufficiently long consecutive sequence of at least

$$\ell = \frac{4(1+3\epsilon)}{\epsilon^2 f [1-(1+\delta)\gamma^2 f]^{\Delta+\Phi+1}} \cdot \max\{\Delta+\Phi,\tau\} \cdot \gamma^3 \cdot \lambda \tag{6}$$

consecutive rounds.

We are now ready to discuss the conditions that protocol parameters should satisfy. We first quantify the length of a clock synchronization interval R, the length of a target recalculation interval M and the length of the convergence phase K. Specifically, we let one target recalculation epoch consists of 4 clock synchronization intervals, i.e., M = 4R; we set $K = \ell + 2\Delta + 4\Phi$ (this will coincide with our common prefix parameter and thus provide some desired properties).

Next, we will require that ℓ (defined in Equation (6)) is appropriately small compared to the length of an epoch and of an interval (note that M = 4R).

$$\ell + 2\Delta + 7\Phi \le \epsilon M/(4\gamma) = \epsilon R/\gamma.$$
(C1)

Further, we require that the advantage of the honest parties is large enough to absorb the errors introduced by ϵ (from the concentration of random variables) and $[1 - (1 + \delta)\gamma^2 f]^{\Delta + \Phi}$ (from the network delay and clock skews).

$$[1 - (1 + \delta)\gamma^2 f]^{\Delta + \Phi} \ge 1 - \epsilon \text{ and } \epsilon \le \delta/12 \le 1/12.$$
 (C2)

4.3 Typical Executions

We define the notion of *typical* executions following [11,12]. The idea here is that given a certain execution E, we compare the actual progress and the expected progress that parties will make under the success probabilities. If the difference and variance are reasonably small, and no bad events (see Definition 6) about the underlying hash function happen, we declare E typical.

Definition 6. An insertion occurs when, given a chain C with two consecutive blocks \mathcal{B} and \mathcal{B}' , a block \mathcal{B}^* created after \mathcal{B}' is such that $\mathcal{B}, \mathcal{B}^*, \mathcal{B}'$ 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.

Note that in addition (compared to [11,12]), in Definition 7(a) we require that the difficulty of all blocks the alert parties can acquire during consecutive rounds S (i.e., D(S)) is well lower-bounded. This is because D(S) also captures the beacon production process, where there is no loss incurred by the boundeddelay network as well as by skewed local clocks. Hence, a reasonably better lower-bound on D(S) helps us get better results when arguing for the good properties of generated beacons by alert parties (Lemma 7).

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

(a) For any set S of at least ℓ consecutive good rounds,

$$(1-\epsilon)[1-(1+\delta)\gamma^2 f]^{\Delta}ph(S) < Q(S) \le D(S) < (1+\epsilon)ph(S)$$

and $D(S) > (1-\epsilon)ph(S).$

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

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

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

In the next lemma, we establish the quantitative relation between honest and adversarial hashing power during consecutive rounds with length at least ℓ , as well as the relationship between the total difficulty acquired by all parties (D(S) + A(J)) and their hashing power.

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

- (a) $(1+\epsilon)p|J| \le Q(S) \le D(U) < (1+5\epsilon)Q(S).$
- (b) $T(J)A(J) < \epsilon M/4(1+\delta)$ or $A(J) < (1+\epsilon)p|J|$; $\tau T(J)B(J) < \epsilon M/4(1+\delta)$ 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 $T(J) \geq T_w^{\min}$, then $A(J) < (1-\delta+3\epsilon)Q(S)$.
- (d) If w is a good round such that $|w r| \leq s$ for any $r \in S$ and $T(J) \geq T_w^{\min}$, then $D(S) + A(J') < (1 + \epsilon)p(h(S) + |J'|)$ where J' denotes the set of adversarial queries in S.

We conclude that almost all executions (that are polynomially bounded by κ and λ) are typical.

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

4.4 Proof Roadmap

In the remainder of this section we present an overview of the analysis. Note that the predicates in Definition 5 are proved in an inductive way over the space of typical executions in a (γ, s) -respecting environment.

First, we focus on the steady block genetation rate. For a warm-up, we argue that an adversarial fork cannot happen too long ago and then extract the common prefix parameter. Equipped with this knowledge, we show that if good skews and certain time adjustment calculations are maintained during a target recalculation epoch, the block production rate will be properly controlled in the next epoch.

Lemma 3. GOODROUND $(r-1) \Longrightarrow$ NoStaleChains(r).

Lemma 4. GOODROUND $(r-1) \land$ GOODSKEW $(r-1) \Longrightarrow$ COMMONPREFIX(r).

Lemma 5. GOODROUND $(r-1) \land$ GOODCHAINS $(r-1) \land$ GOODSKEW $(r-1) \land$ GOODSHIFT $(r-1) \Longrightarrow$ BLOCKLENGTH(r).

Lemma 6. $\operatorname{GOODROUND}(r-1) \Longrightarrow \operatorname{GOODCHAINS}(r)$.

Corollary 1. GOODROUND $(r-1) \Longrightarrow$ GOODROUND(r).

Next, we move to the properties w.r.t. clocks. We argue that if at the onset, the PoW difficult is appropriately set and the steady block generating rate lasts during the whole clock synchronization interval, the beacon set used by parties to update their clock will be identical and the majority of these beacons will be produced and emitted by alert parties. For synchronized parties, this good beacon set implies that the differences between alert parties' local clocks are still narrow after they enter the next interval and that the shift value they computed is well-bounded. Furthermore, regarding newly joining parties, we also provide an analysis of the joining procedure showing that joining parties starting with no *a-priori* knowledge of the global time, they can listen in and bootstrap their logical clock and become alert parties. The above two aspects imply that a bounded skew is maintained over the whole execution.

Lemma 7. $\operatorname{GOODROUND}(r-1) \Longrightarrow \operatorname{GOODBEACONS}(r)$.

Lemma 8. $\text{GOODSKEW}(r-1) \Longrightarrow \text{GOODSHIFT}(r)$

Lemma 9. GOODSKEW $(r-1) \land$ GOODBEACONS $(r-1) \Longrightarrow$ GOODSKEW(r).

To sum up, a "safe" start and a (γ, s) -respecting environment guarantee that good properties can be achieved during the whole execution.

Theorem 2. For a typical execution in a $(\gamma, M + 2(\ell + 2\Delta + 7\Phi))$ -respecting environment, if Condition C1 and Condition C2 are satisfied, then all predicates in Definition 5 hold.

Finally, we work out the related parameters (in Theorem 3) and conclude that Timekeeper solves the clock synchronization problem.

Theorem 3. Consider an execution of Timekeeper in a $(\gamma, M+2(\ell+2\Delta+7\Phi))$ respecting environment. If Conditions (C1) and (C2) are satisfied, then the protocol achieves clock synchronization (Definition 1) with parameter values

Skew = 2Φ , shiftLB = $3\Phi/R$, shiftUB = $2\Phi/R$,

except with probability negligibly small in κ and λ .

References

- Abraham, I., Devadas, S., Dolev, D., Nayak, K., Ren, L.: Synchronous byzantine agreement with expected O(1) rounds, expected o(n²) communication, and optimal resilience. In: Goldberg, I., Moore, T. (eds.) Financial Cryptography and Data Security - 23rd International Conference, FC 2019, Frigate Bay, St. Kitts and Nevis, February 18-22, 2019, Revised Selected Papers. Lecture Notes in Computer Science, vol. 11598, pp. 320–334. Springer (2019). https://doi.org/10.1007/ 978-3-030-32101-7_20, https://doi.org/10.1007/978-3-030-32101-7_20
- Badertscher, C., Gaži, P., Kiayias, A., Russell, A., Zikas, V.: Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security. pp. 913–930. CCS '18, Association for Computing Machinery, New York, NY, USA (2018). https://doi.org/10.1145/3243734.3243848
- Badertscher, C., Gaži, P., Kiayias, A., Russell, A., Zikas, V.: Dynamic ad hoc clock synchronization. In: Canteaut, A., Standaert, F.X. (eds.) Advances in Cryptology – EUROCRYPT 2021. pp. 399–428. Springer International Publishing, Cham (2021)
- Badertscher, C., Maurer, U., Tschudi, D., Zikas, V.: Bitcoin as a transaction ledger: A composable treatment. In: Katz, J., Shacham, H. (eds.) Advances in Cryptology - CRYPTO 2017. pp. 324–356. Springer International Publishing, Cham (2017)
- Bagaria, V., Kannan, S., Tse, D., Fanti, G., Viswanath, P.: Prism: Deconstructing the blockchain to approach physical limits. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security. pp. 585–602. CCS '19, Association for Computing Machinery, New York, NY, USA (2019). https://doi.org/10.1145/3319535.3363213
- Bahack, L.: Theoretical bitcoin attacks with less than half of the computational power (draft). Cryptology ePrint Archive, Report 2013/868 (2013), https://ia. cr/2013/868

- 30 Juan Garay, Aggelos Kiayias, and Yu Shen
- Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: 42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA. pp. 136–145. IEEE Computer Society (2001). https://doi.org/10.1109/SFCS.2001.959888
- Dolev, D., Halpern, J.Y., Strong, H.R.: On the possibility and impossibility of achieving clock synchronization. J. Comput. Syst. Sci. 32(2), 230-250 (1986). https://doi.org/10.1016/0022-0000(86)90028-0
- Dwork, C., Lynch, N., Stockmeyer, L.: Consensus in the presence of partial synchrony. J. ACM 35(2), 288-323 (Apr 1988). https://doi.org/10.1145/42282. 42283
- Garay, J., Kiayias, A., Leonardos, N.: The bitcoin backbone protocol: Analysis and applications. In: Oswald, E., Fischlin, M. (eds.) Advances in Cryptology -EUROCRYPT 2015. pp. 281–310. Springer Berlin Heidelberg, Berlin, Heidelberg (2015)
- Garay, J., Kiayias, A., Leonardos, N.: The bitcoin backbone protocol with chains of variable difficulty. In: Katz, J., Shacham, H. (eds.) Advances in Cryptology – CRYPTO 2017. pp. 291–323. Springer International Publishing, Cham (2017)
- Garay, J., Kiayias, A., Leonardos, N.: Full analysis of nakamoto consensus in bounded-delay networks. Cryptology ePrint Archive, Report 2020/277 (2020), https://ia.cr/2020/277
- Garay, J., Kiayias, A., Ostrovsky, R.M., Panagiotakos, G., Zikas, V.: Resourcerestricted cryptography: Revisiting mpc bounds in the proof-of-work era. In: Canteaut, A., Ishai, Y. (eds.) Advances in Cryptology – EUROCRYPT 2020. pp. 129– 158. Springer International Publishing, Cham (2020)
- Garay, J., Kiayias, A., Shen, Y.: Permissionless clock synchronization with public setup. Cryptology ePrint Archive, Report 2022/1220 (2022), https://eprint. iacr.org/2022/1220
- 15. Garay, J.A., Kiayias, A.: Sok: A consensus taxonomy in the blockchain era. In: Jarecki, S. (ed.) Topics in Cryptology - CT-RSA 2020 - The Cryptographers' Track at the RSA Conference 2020, San Francisco, CA, USA, February 24-28, 2020, Proceedings. Lecture Notes in Computer Science, vol. 12006, pp. 284–318. Springer (2020). https://doi.org/10.1007/978-3-030-40186-3_13
- Halpern, J.Y., Simons, B., Strong, R., Dolev, D.: Fault-tolerant clock synchronization. In: Proceedings of the Third Annual ACM Symposium on Principles of Distributed Computing. pp. 89–102. PODC '84, Association for Computing Machinery, New York, NY, USA (1984). https://doi.org/10.1145/800222.806739
- Katz, J., Maurer, U., Tackmann, B., Zikas, V.: Universally composable synchronous computation. In: Proceedings of the 10th Theory of Cryptography Conference on Theory of Cryptography. pp. 477–498. TCC'13, Springer-Verlag, Berlin, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36594-2_27
- Lamport, L.: Time, clocks, and the ordering of events in a distributed system. Commun. ACM 21(7), 558-565 (1978). https://doi.org/10.1145/359545.359563
- Lamport, L., Melliar-Smith, P.M.: Byzantine clock synchronization. In: Proceedings of the Third Annual ACM Symposium on Principles of Distributed Computing. pp. 68–74. PODC '84, Association for Computing Machinery, New York, NY, USA (1984). https://doi.org/10.1145/800222.806737
- Lenzen, C., Loss, J.: Optimal clock synchronization with signatures. In: Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing. pp. 440–449. PODC'22, Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3519270.3538444

31

- 21. Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system (2008), http://bitcoin.org/bitcoin.pdf
- 22. Pass, R., Seeman, L., Shelat, A.: Analysis of the blockchain protocol in asynchronous networks. In: Coron, J., Nielsen, J.B. (eds.) Advances in Cryptology EUROCRYPT 2017 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 May 4, 2017, Proceedings, Part II. Lecture Notes in Computer Science, vol. 10211, pp. 643–673 (2017). https://doi.org/10.1007/978-3-319-56614-6_22
- Pass, R., Shi, E.: Fruitchains: A fair blockchain. In: Proceedings of the ACM Symposium on Principles of Distributed Computing. pp. 315–324. PODC '17, Association for Computing Machinery, New York, NY, USA (2017). https://doi.org/10.1145/3087801.3087809
- Pass, R., Shi, E.: Rethinking large-scale consensus. In: 30th IEEE Computer Security Foundations Symposium, CSF 2017, Santa Barbara, CA, USA, August 21-25, 2017. pp. 115–129. IEEE Computer Society (2017). https://doi.org/10.1109/CSF.2017.37
- Srikanth, T.K., Toueg, S.: Optimal clock synchronization. J. ACM 34(3), 626–645 (1987). https://doi.org/10.1145/28869.28876
- Welch, J.L., Lynch, N.: A new fault-tolerant algorithm for clock synchronization. Information and Computation 77(1), 1–36 (1988)

A Glossary

Parameter	Description	
h_r	The number of honest RO queries in (nominal) round r .	
t_r	The number of RO quries by \mathcal{A} in (nominal) round r .	
δ	Advantage of honest parties $(t_r < (1 - \delta)h_r$ for all $r)$.	
f	The probability at least one honest RO query out of n_0 computes a	
	block for target T_0 .	
R	The length of a clock synchronization interval in number of rounds.	
М	The length of a target recalculation epoch in number of rounds.	
К	The length of convergence phase in a clock synchronization interval in	
	number of rounds.	
Δ	Network delay in rounds.	
$\Phi_{ m clock}$	The upper bound of the drift that \mathcal{A} can set.	
Φ	The upper bound of the difference between honest parties' local clocks.	
	We require that $\Phi = \Phi_{\text{clock}} + \Delta$.	
κ	Security parameter; length of the hash function output.	
(γ, s)	Respecting environment parameter.	
ϵ	Quality of concentration of random variables.	
λ	Related to the properties of the protocol.	

Table 3. Main Parameters of Timekeeper.