

# Ouroboros Chronos: Permissionless Clock Synchronization via Proof-of-Stake

Christian Badertscher<sup>1\*</sup> , Peter Gazi<sup>1</sup>, Aggelos Kiayias<sup>1,2\*\*</sup>, Alexander Russell<sup>1,3\*\*\*</sup> , and Vassilis Zikas<sup>2†</sup>

<sup>1</sup> IOHK – `firstname.lastname@iohk.io`

<sup>2</sup> University of Edinburgh – `firstname.lastname@ed.ac.uk`

<sup>3</sup> University of Connecticut – `acr@cse.uconn.edu`

**Abstract.** Clock synchronization allows parties to establish a common notion of global time by leveraging a weaker synchrony assumption, i.e., local clocks with approximately the same speed. The problem has long been a prominent goal for fault-tolerant distributed computing with a number of ingenious solutions in various settings. However, despite intensive investigation, the existing solutions do not apply to common blockchain protocols, which are designed to tolerate variable—and potentially adversarial—participation patterns, e.g., sleepiness and dynamic availability. Furthermore, because such blockchain protocols rely on freshly joining (or re-joining) parties to have a common notion of time, e.g., a global clock which allows knowledge of the current protocol round, it is not clear if or how they can operate without such a strong synchrony assumption.

In this work, we show how to solve the global synchronization problem by leveraging proof of stake (PoS). Concretely, we design and analyze a PoS blockchain protocol in the above dynamic-participation setting, that does not require a global clock but merely assumes that parties have local clocks advancing at approximately the same speed. Central to our construction is a novel synchronization mechanism that can be thought as the blockchain-era analogue of classical synchronizers: It enables joining parties—even if upon joining their local time is off by an arbitrary amount—to quickly calibrate their local clocks so that they all show approximately the same time. As a direct implication of our blockchain construction—since the blockchain can be joined and observed by any interested party—we obtain a permissionless PoS implementation of a global clock that may be used by higher level protocols that need access to global time.

## 1 Introduction

Global clock synchronization [15,27,20] allows a set of mutually distrustful parties to approximate on a global notion of “time,” in such a manner that if some party believes that the global time is  $\tau$  then every party believes it to be  $\tau \pm \epsilon$  for some small  $\epsilon > 0$ . This in terms, allows for an (approximately) synchronous (or partially synchronous) execution of distributed protocols which has placed the study of such *synchronizers* at a prominent position in the theoretical computer science research. A number of works investigating feasibility across the spectrum of security/adversary models—from perfect to computational security and for different types of network synchronisation assumption [15,27,20,17,16,2,37,28,32,36]. (For completeness, we include a description of the current landscape of feasibility in Appendix A.) The common assumption of such synchronizers is that the (honest) parties have local (initially desynchronized) clocks which advance at (roughly) the same speed.

Notwithstanding, existing synchronization techniques are inapplicable to the standard model used for the analysis of Nakamoto-style blockchain/consensus protocols. The reason is that the above traditional models assume accurate knowledge of the total number of parties present in the system, and their main

---

\* Work done while the author was at the University of Edinburgh, Scotland.

\*\* Research partly supported by EU Project No. 780477, PRIVILEGE.

\*\*\* This material is based upon work supported by the National Science Foundation under Grant No. 1717432.

† Work supported in part by IOHK and done in part while the author was visiting the Simons Institute for the Theory of Computing, UC Berkeley.

tool is smart counting of messages (or message chains) received—a standard technique in the distributed computing literature that is also adopted in “iterated” Byzantine Fault Tolerant (iBFT) protocols. In contrast, in the Nakamoto-style setting, Pass and Shi put forth the *sleepy model* of consensus as a natural model for the analysis of such protocols, which reduces the applicability of such counting-based approaches. The reason is that the sleepy model allows for parties to join the protocol at any time and then to temporarily *sleep*—i.e., drop out of the protocol<sup>4</sup>—according to an arbitrary (or even adversarial) sleep pattern. This model was later generalized—and adapted to the (G)UC setting [9,10]—in the *dynamic availability* model of [4], which captures sleepiness with respect to arbitrary resources available to the protocol, e.g, the clock (capturing temporary loss of time), the network (capturing temporary connectivity issues), or the random oracle (temporary unavailability of computing resources).

The above dynamic participation assumptions, however natural, limit the power of existing synchronization techniques, since the lack of agreement of participation patterns makes counting ineffective to take consistent decisions without additional synchrony assumptions. In fact, the cryptographic analysis of formally specified and analyzed proof-of-work (PoW) and proof-of-stake (PoS) blockchains, has typically assumed a (partially) synchronous model with a notion of global time. Indeed, standard references for the proven security of Bitcoin [18,19,33] implicitly use the fact that they can refer to a global round index in order to prove the desired properties of the protocol. For example, the common-prefix property is defined to require that if an honest party holds a chain at Round  $\rho$ , then the prefix of this chain—obtained by removing the  $k$  most recent blocks—will eventually become prefix of the chain of any honest party (at some round  $\rho' \geq \rho$ ). The assumption was made explicit in [5] by assuming a global clock in the global UC setting [10]: this permits every party to query a common clock on demand and from that deduce the current round. A similar approach, assuming access to a global clock, was also adopted in the constructions of PoS blockchains, such as Sleepy Consensus [35], Snow White [6], and Ouroboros [26,14,4]. So a natural question arises: Do the above blockchain protocols preserve their security when the (implicit) agreement on the current round<sup>5</sup> is replaced by the assumption of local potentially unsynchronized clocks that proceed at roughly the same speed?

Interestingly, answering this question reveals a delicate distinction between proof-of-work (PoW) based protocols, in particular Bitcoin, and common Nakamoto-style proof-of-stake (PoS) based protocols. In particular, the description of the Bitcoin blockchain (without difficulty recalibration) can rely on a purely *execution-driven* notion of time: when a party engages in the protocol it need not be aware of the current global time. In fact, in the static difficulty setting, proving security in this way follows immediately from [18,33]. Moreover, a notion of global time may be inferred by the current blockchain length in each party’s local state. In the PoS setting, however, this execution-driven notion of time has been achieved only by Algorand [12], requiring a highly relevant concession: that explicit participation bounds are known to the protocol participants, i.e., each protocol participant at any given time is aware of how many protocol parties are expected to engage in the protocol step. Such a rigid participation restriction is not necessary for the Bitcoin blockchain.

This state of affairs leaves an important gap between the PoW setting and those PoS protocols that require no explicit participation bounds such as [35,6,26,14,4]: these protocols use a notion of global time hardwired in the protocol logic, and the protocol is unspecified without such global knowledge of time/round. In fact, there does not seem to be a simple way to removing this dependence of global time, and replace it by local clocks (even, perfectly-coordination ones that advance at exactly the same speed) while preserving the security guarantees. (Of course, one could include such a notion of (approximate) global time in a trusted *checkpointing assumption* [13], but this defeats the purpose of decoupling the protocol from an explicitly assumed trusted source of global time when joining the protocol which, as discussed below, is the main challenge of our work.)

In a nutshell, the reason is that in these PoS blockchains a party’s right to create a block is always associated with a concrete round<sup>6</sup> (also called “slot”), and in order to verify that a block is created by an eligible party, that party must include a proof explicitly referring to the slot number. This means that a

<sup>4</sup> Different degrees of sleepiness allow for queueing received messages which are read once the sleepy party re-joins (light sleepers) or even missing the messages sent while the party is asleep (heavy sleepers)

<sup>5</sup> Recall that this implicit agreement was abstracted as an explicit global clock in [5,4].

<sup>6</sup> One can view this as an implicit timestamp

new party that joins the blockchain—or one that has been sleeping for long—cannot prune-off chains with adversarial timestamps so that he eventually adopts the right chain. Thus if a new party with an incorrect local time joins the protocol and sees a chain that includes blocks which appear to be far in the future (according to her local time), she cannot decide whether the chain is adversarial—in which case she needs to ignore or truncate it—or her local time is far behind absolute time. It is worth adding that these are not merely theoretical considerations: in a real world deployment the dependency on a global clock is typically met by using a global time synchronization service, e.g., NTP [30] and hence the security of all these protocols becomes compromised if such service fails to deliver a truly reliable clock. This is a possibility that cannot be excluded [29].

Note that all previous PoS protocols which can operate in a participation-unrestricted setting [35,6,26,14,4] require an upper bound on the network delay  $\Delta$  to be known to all participants; this in fact is a necessary assumption, see [34], due to the participation uncertainty. Unfortunately, knowing an upper bound on  $\Delta$  does not help the parties in any direct way to assess the actual time (e.g., by locally counting time intervals of length  $\Delta$ ), as participation gaps can invalidate their local timer with respect to the implicit global execution-driven time.

**Our contributions.** We present the first provably secure approach to global clock synchronization in the dynamic participation setting from standard Nakamoto-style consensus assumptions. We focus on PoS—where we assume the same honest stake-majority underlying existing PoS protocols [4]—but we expect that our techniques can be adapted to the PoW setting. To this direction, we devise the first participation-unrestricted (i.e., without explicitly known participation bounds; see below) PoS blockchain that does not need a global clock and, instead, enables parties to implement and maintain a bounded-drift clock in the delayed-delivery network model [33,19]. Our new blockchain protocol avoids the dependency on an external service providing access to global time, which constitutes an improvement in resiliency compared to previous works. Additionally, our protocol can be used as a cryptographically secure synchronizer by any external application in this new era of dynamic participation.

In more detail, our construction assumes that the inaugural parties, i.e., those that initially commence the protocol execution (but might later-on sleep or drop out), have access to local clocks indicating their local times which might be off by a parameter  $\Delta$  and which advance at (roughly) the same speed. This is similar in spirit to the guarantee offered by the timing model of Kalai et al. [22]; however, to guarantee full compatibility with past (G)UC statements—including universal composability—we introduce a new, imperfect version of the clock functionality [24] as a global setup, which captures that parties advance at roughly the same pace but, unlike previous attempts to capture synchrony in a composable framework [31,21,3], allows parties to advance to a next round even before every honest party has finished with his current round, as long as front-running parties do not drift too far ahead. We believe that this global-UC formulation of relaxed synchronization can be of independent interest.<sup>7</sup>

Our protocol guarantees that parties who later join the protocol, no matter how outdated their local clock value is, will synchronize themselves (up to a small drift depending on network delay) with the rest of the parties, and remain synchronized for as long as they faithfully execute the protocol.

Although we conduct our formal analysis in the UC framework (where we model certain setup functionalities as being globally accessible) [9,10], we also provide informal description and analysis of our protocols for readers not familiar with the framework. Note that the advantage of UC analysis is that it ensures that the resulting synchronizer can be used to enable the design of synchronous protocols—which enjoy the benefits sketched above—from our arguably weak assumptions of similar speed local clocks and a bounded-delivery network. In fact, as the protocol exports the clock it implements to all participants, a party wishing to only use it for synchronization (for a separate application) can simply join the protocol *in a passive mode*—without any need to have stake or participate in the lottery as a slot leader—and keep executing it for a sufficiently large number of (locally clocked) steps. By doing so, the party will compute a clock that, as we prove, is

---

<sup>7</sup> We note that there are UC realizations of other “clocks” in the literature, such as [11]. This is an unfortunate name-clash of different concepts: That clock is a much weaker functionality that does not guarantee even a relaxed (bounded-drift) round-based structure with guaranteed termination.

guaranteed to be in (loose, up to a small constant offset) synchronization with all honest parties running the protocol (in a passive or full-participation mode) as long as (available) honest stake majority is preserved.

Our construction builds on Ouroboros Genesis [4], where to signify the connection, we refer to our protocol as (*Ouroboros*) *Chronos*. Similar to [4] we work in the dynamic-availability framework which implies that our solution works independently of the exact number of honest parties who are around, as long as they hold a majority of the system’s (currently available) stake. The design and analysis of Chronos uses, in a white-box manner, elements from previous PoS protocols [35,6,26,14,4], but involves several new ideas which are crucial in achieving global synchronization in this participation-unrestricted setting. We discuss these ideas in detail below, where we give an overview of our design. (We include details and points of comparison especially with the most recent prior work [4] in Section C.)

**Overview of our techniques.** The heart of our protocol, namely its procedure for synchronization of newly joining parties, starts with these parties listening on the network for some time, collecting broadcasted chains and following a “densest chain” chain-selection rule. Informally, this rule mandates that if two chains  $\mathcal{C}$  and  $\mathcal{C}'$  start diverging at some time  $\tau$ —according to the reported time-stamps in  $\mathcal{C}$  and  $\mathcal{C}'$ —then prefer the chain which is denser in a sufficiently long interval after that time. Our first key observation is that this rule offers a useful (albeit in itself insufficient) guarantee in our setting: the joining party will end up with some blockchain that, although *arbitrarily long*, is at worst forking from a chain held by an honest and already synchronized party by a bounded number of blocks (equal to a security parameter) with overwhelming probability. This observation is the key to start building our synchronization mechanism. More concretely, we prove that the above process guarantees to eventually prune-off all chains with bad prefixes, i.e., prefixes that do not largely coincide with the prefixes of the other already synchronized honest parties’ chains. In fact, as we show, the parties can compute an upper bound on the time (according to their local clocks) they need to remain in the above self-synchronization state before they build confidence in the above guarantee, i.e., before they know that their locally held chain is consistent with a long and stable prefix that already-synchronized honest parties adopt.

The second key observation is that once a joining party has converged to such a *fresh*—i.e., produced after the joining party was activated—prefix of an honest chain, it may use the difference between its current local time and the (local) time recorded when this chain (and other control information) was received to adjust its local clock so that its local time is consistent with the times reported on the prefix. The hope would be that a clever adjustment will bring this local time sufficiently close to that of an honest and already synchronized party.

Designing and analyzing such a natural updating process is unexpectedly challenging. To see why, consider the following naïve attempt: The party resets its local clock so that the time reported in, say, the last block of the prefix is the time this block was received. Before discussing the limitations of this proposal, let us first discuss an inherent property when dealing with clock synchronization in the setting with  $\Delta$ -bounded (but adversarially controlled) delay networks. A message received by a party might have been sent up to  $\Delta$  rounds before, hence the time that the party will set its clock to might be up to  $\Delta$  rounds away from the clock of the sender (at the point of update). This delay-induced imprecision is unavoidable, so when we assess a given proposal we accept that clocks only need to be “loosely” synchronized; specifically, clocks of honest parties might differ by a bounded amount, where the bound is known and depends only on  $\Delta$ . In fact, this relaxation is common and believed to be necessary even in the permissioned model [27,20].<sup>8</sup>

However, the above simple solution is problematic even when no delays are there: Although the chain that the newly joining party recovered is guaranteed to have a prefix consistent with the already synchronized honest parties, individual blocks might be originating from the adversary and therefore contain a time stamp very different from the true sending time of that block. To make matters worse, the rate of honestly generated blocks in a chain of an honest party can be quite low as implied by the known bounds of chain quality [19,14], and thus the time inaccuracy of any individual block can be significant.

A second attempt would be to have in every round (or at regular intervals) every party use the credentials of all the coins it owns to broadcast a signed timestamp, i.e., every party acts as a verifiable *synchronization*

<sup>8</sup> The model from [27] with honest clocks that report values differing by up to  $\Delta$  is equivalent to a situation in which clocks report the right value, but parties might receive it with a difference of up to  $\Delta$  rounds.

(or *timestamping*) *beacon* on behalf of all the coins it owns. The joining party receives all these broadcasted timestamps, and uses their majority to compute the value of its clock. Still, this solution has drawbacks: The first is scalability; this is not severe, as existing ideas can be employed such as using the protocol history as input to a verifiable random function (VRF) to identify eligible parties (or, as in the case of Algorand, by using Bracha-style committees [7]) to send timestamping beacons in every synchronization round. The second, harder problem is that in order to use the majority, the local clocks of the parties that report time need to be perfectly synchronized so that their majority agrees. If their clocks have any small drift, this fails. Furthermore, even with identical speed clocks, dynamic participation allows parties to drop off and rejoin, which means that, due to the network delay the honest parties will end up with only loosely synchronized local clocks. Using the average instead of the majority function does not help out here either since a single adversarial timestamp can throw off the average arbitrarily far. Hence, taking the median of the received timestamps promises to be more stable against extreme values. Observe that as long as synchronized honest parties' local clocks are not far apart, the times they report will be concentrated to a sufficiently small time interval, and the median will fall in this interval.

The above insight brings us closer, but is still insufficient: If the adversary can serve to, say, two different joining parties different and possibly disjoint sets of timestamps (on behalf of eligible corrupted synchronization-beacon parties) then he could force an opposing clock adjustment between the two that will increase their clock drift well beyond the drift of any pair of already synchronized parties. To resolve this issue, we need to ensure that the parties agree on the set of eligible timestamps (whether honest or corrupted) that they use for adjusting their local time. This is a classical consensus problem. Luckily, our synchronizer runs in tandem with a PoS-based blockchain which solves consensus with dynamic availability, and which can assist in reaching agreement on the synchronization-beacon values for recalibration. And thanks to the property discussed at the beginning of the section—namely that even joining parties (without accurate time) will eventually be able to bootstrap a sufficiently long prefix of the blockchain—the joining parties will agree on the set of beacons for recalibration.

Our technical solution follows the spirit of the final conclusion above. In a nutshell, we will use the VRF to assign timestamping-beacon parties to slots according to their state. Parties who are synchronized and active when their assigned slot is encountered will broadcast a timestamp and a VRF-proof of their eligibility for the current timeslot (together, we call this a *synchronization beacon*). And to agree on the set of eligible parties that will be used (including the dishonest ones) these beacons will also be included in the blockchain by the already synchronized parties, similarly to transactions. Any party who joins and tries to get synchronized will gather chains and record any broadcasted beacons (and keep track of the local time these were received). Once the party is confident it has a sufficiently long prefix of the honest chain, it will retrospectively use this gathered information to extract the agreed-upon set of beacons, compute a good approximation of the clocks parties had when they broadcasted these beacons and apply a median rule to set its local clock to at most a small distance from other honest and synchronized parties. In order to ensure that already synchronized parties adjust in tandem with joining parties we will have them also periodically execute the synchronization algorithm—but of course using their local blockchain, which they know is guaranteed to have a large common prefix with any other honest and synchronized party. Evidently, to turn this high-level idea of our solution into a provably secure protocol requires careful design choices that we present in Section 3. The analysis is given in Section 6. This constitutes the main technical contribution of this paper.

**Outline.** Section 2 sketches our model and Section 3 describes our main result, i.e., outlines the Chronos protocol and the functionality it realizes, along with the security analysis. These sections are written in a self-contained manner so that the reader can get a first idea of our contributions and techniques. Sections 4-6 include more details on our results and can help an interested reader to get an overview of the detailed treatment that is presented in the Appendix.

## 2 Our Model

**Basic Notation.** For  $n \in \mathbb{N}$  we use the notation  $[n]$  to refer to the set  $\{1, \dots, n\}$ . For brevity, we often write  $\{x_i\}_{i=1}^n$  and  $(x_i)_{i=1}^n$  to denote the set  $\{x_1, \dots, x_n\}$  and the tuple  $(x_1, \dots, x_n)$ , respectively. For a tuple

$(x_i)_{i=1}^n$ , we denote by  $\text{med}((x_i)_{i=1}^n)$  the (lower) median of the tuple, i.e.,  $\text{med}((x_i)_{i=1}^n) \triangleq x'_{\lceil n/2 \rceil}$ , where  $(x'_i)_{i=1}^n$  is a (non-decreasing) sorted permutation of  $(x_i)_{i=1}^n$ .

For a blockchain (or chain)  $\mathcal{C}$ , which is a sequence of blocks, we denote by  $\mathcal{C}^{\lceil k}$  the chain that is obtained by removing the last  $k$  blocks; and by  $\text{head}(\mathcal{C})$  the last block of  $\mathcal{C}$ . We write  $\mathcal{C}_1 \preceq \mathcal{C}_2$  if  $\mathcal{C}_1$  is a prefix of  $\mathcal{C}_2$ .

**Dynamic Availability.** We adopt the dynamic availability framework from [4] which captures parties joining and leaving the protocol at (the environment’s) will. This is done by equipping the functionalities, global setups, and the protocol with explicit registration/de-registration commands, thereby keeping track of when parties are joining and adjusting their guarantees depending based on this information. We refer the reader to [4] for a detailed discussion of this model.

**Relaxed Synchrony.** The synchrony assumption that parties advance at exactly the same pace can be captured by the global-setup variant of the clock functionality from [24]. This is a weaker version of the global clock used in previous composable analyses of blockchains [5,4] in that it does not keep a counter representing the global system time, but rather maintains for each party (resp. ideal functionality) an indicator-bit  $d_{\mathcal{P}}$  (resp.  $d_{(\mathcal{F}, \text{sid})}$ ) of whether or not a new round has started. Each party’s indicator is accessible by a standard CLOCK-GET command. All indicators are set to 0 at the beginning of each round; once any party or functionality finishes its round it issues a CLOCK-UPDATE command that updates his indicator to 1. Once every party and functionality has updated its indicator, the clock resets all of them to 0; this switch allows the parties to detect that the previous round has ended and move on to the next round.

Arguably the above clock offers very strong synchronisation guarantees, since once a round switches, every party is informed about it in the next activation. In [24] a relaxed version of this clock was introduced which allowed the adversary to delay notifying the parties about a round switch by bounded amount of fetch-attempts. This behavior relaxes the perfect nature of the clock, but it still ensures that no party advances to a next round before all parties have completed their current round.

In this work we consider parties that advance at roughly the same speed, which means that a party might advance its round even before another party has finished with its current round, and even multiple times, as long as its is ensured that no honest party is left too far behind. For this purpose we introduce an even more relaxed version of the (global-setup variant) of the clock from [24] which, intuitively, allows a party to advance to its next round multiple times *before* some honest parties have completed their current round, as long as the relative pace of advancement for any two honest parties stays below a drift parameter  $\Delta_{\text{clock}}$ . We note in passing that a similar guarantee was formulated in the timing model [23]; however, the solution there notified the underlying model of computation which creates complications with the (G)UC composition theorem which would need to be reproved. To avoid such complications, in this work we capture the above relaxed synchrony assumption as a global functionality.<sup>9</sup>

The above is captured as follows: Similar to the perfect clock above, the imperfect clock stores an indicator-bit  $d_{\mathcal{P}}$  which is used to keep track of when everyone has completed a round (not necessarily the same round)—one can think of this indicator as corresponding to a baseline round-switch, which is however hidden from the parties and might only be observed by ideal functionalities. Additionally, for every party the imperfect clock keeps an imperfect version of the indicator bit  $d_{\mathcal{P}}^{\text{Imp}}$  (corresponding to switches  $\mathcal{P}$ ’s *local*, e.g., hardware, clock switches) which is what is exported when the party attempts to check his clock.

This local indicator is used similarly to how synchronous protocols would use the perfect indicator in [24]; but we allow the adversary to control when this local indicator is updated under the restrictions that (a)  $d_{\mathcal{P}}^{\text{Imp}}$  cannot advance in the middle of  $\mathcal{P}$ ’s round, (b) it cannot fall behind the baseline induced by the indicator  $d_{\mathcal{P}}$ , and (c) it cannot advance ahead of the baseline by more than  $\Delta_{\text{clock}}$ . This is achieved by the imperfect clock keeping track of the relative difference/distance  $\text{drift}_{\mathcal{P}}$  between the number of local advances of each registered  $\mathcal{P}$  from the baseline updates; this distance is increased whenever  $d_{\mathcal{P}}^{\text{Imp}}$  is reset (by the adversary) to 0 and decreased whenever the baseline indicator  $d_{\mathcal{P}} \in \{0, 1\}$  is reset to 0; if the distance of some party

<sup>9</sup> In [24] a functionality corresponding to the timing-model assumptions [23] was proposed along with a reduction to the (local) clock functionality. However, both the fact that their clock functionality is local and that their reduction uses a complete network of (known) bounded-delay authenticated channels—which we do not assume here—makes that result incompatible with our model and goals.

from the baseline falls below 0 (i.e., the adversary attempts to stall a party when the baseline advances<sup>10</sup>) then the local indicator is reset to  $d_P^{Imp} = 0$  (which allows P to advance his round) and the corresponding distance is also reset to 0.

**Modeling Peer-to-Peer Communication.** We assume a diffusion network in which all messages sent by honest parties are guaranteed to be fetched by protocol participants after a specific delay  $\Delta_{\text{net}}$ . Additionally, the network guarantees that once a message has been fetched by an honest party, this message is fetched by any other honest party within a delay of at most  $\Delta_{\text{net}}$ , even if the sender of the message is corrupted. We note that this network model is not substantially stronger than in previous works [5,4], which use a network functionality providing bounded-delay message delivery. Our model is equivalent via an unconditional reduction: echoing received messages. In practice, this reduction of course needs to be applied prudently to avoid saturating the network. This is exactly done by the relevant networking protocols: e.g. in Bitcoin, when a new block is received its hash is advertised and then propagated and validated by the network as needed. Chronos can use the same mechanism. We detail the corresponding functionality in Section B.

**Genesis Block Distribution; the Weak Start Agreement.** Our model allows parties’ local time-stamps to drift apart over the course of an execution; additionally the model makes no assumption that the initialization of the initial stakeholders is completed in the same round, i.e., honest parties might start producing blocks for logical slot 1 in different rounds of the (global) execution. To this aim, we weaken the functionality  $\mathcal{F}_{\text{INIT}}$  adopted by [4] to allow for bounded delays when initial stakeholders receive the genesis blocks. Namely, our  $\mathcal{F}_{\text{INIT}}^{\Delta_{\text{net}}}$  functionality merely guarantees genesis block delivery to initial stakeholder not more than  $\Delta_{\text{net}}$  rounds apart from each other; the offsets are under adversarial control. The details of the  $\mathcal{F}_{\text{INIT}}^{\Delta_{\text{net}}}$  functionality appear in Section B.

**Further Hybrids.** The protocol makes use of a VRF (verifiable random function) functionality  $\mathcal{F}_{\text{VRF}}$ , a KES (key-evolving signature) functionality  $\mathcal{F}_{\text{KES}}$ , and a (global) random oracle functionality  $\mathcal{G}_{\text{RO}}$  (to model ideal hash functions). Full model details are provided in Section B for reference.

### 3 The Synchronizing Blockchain

In this section we provide our main result: a concrete blockchain protocol, called *Chronos* (meaning “time” in greek), implementing a new functionality that extends the ledger from [4] with a robust notion of time. Informally, in addition to the standard ledger guarantees, this new ledger exports a notion of global time with strong accuracy guarantees.

In a nutshell, our new “timed” ledger functionality maintains an immutable ledger state denoted by `state` which encodes a sequence of transaction blocks. How fast this state grows and which transactions it includes are specified in a ledger policy description, which is a ledger parameter. Each registered party can request to see the state, and is guaranteed to receive a sufficiently long prefix of it—the length of each party’s visible prefix of the state at any given point is captured by (monotonically) increasing pointers that point to the last block this party’s current prefix includes. The adversary is given policy-limited control to set these pointers within a sliding window of width `windowSize`, which starts at the head of the state and slides as the state increases. The above mechanism captures the guarantees of the *common-prefix (CP)* property of blockchains (where the common-prefix parameter relates to the width of the window).

Parties advance the ledger when they are instructed to (activated with specific maintain-ledger input) by their UC environment. The ledger uses these queries plus the knowledge of the baseline speed to ensure the advancement of the state (which encapsulates a property analogous to *chain growth (CG)* [18,33]). Per the ledger policy, any party can input a transaction to the ledger, which are validated and guaranteed to be included eventually whenever they stay valid (this will ensure transaction liveness [18]). The ledger gives different guarantees to honest parties depending on whether they are considered synchronized or not. Roughly speaking, synchronized parties are the parties for which one can provide the best security guarantees (for example, they are inaugural, or part of the system for a long time and correctly bootstrapped).

<sup>10</sup> Note that, by definition the baseline advances when all parties have completed their current round.

De-synchronized honest parties receive reduced security guarantees and could, until they become synchronized, have unreliable values.

In our ledger description, the above policy is specified by means of a ledger-parameter/procedure which is denoted as `ExtendPolicy`. The basic mode of operation of `ExtendPolicy` is as follows: it takes, as an input, a proposal from the adversary for extending the state, i.e., a sequence of new blocks to be added, and can then decide to follow this proposal if it satisfies its policy sketched above; if it does not, `ExtendPolicy` can ignore the proposal (and enforce a default extension). It will enforce minimal chain growth, a certain fraction of “good blocks,” (a.k.a. *chain quality* (CQ) [18,33]) and transaction liveness guarantees for old and (still) valid transactions. In fact, as discussed below, in the security proof of Chronos, we will demonstrate that it satisfies the above properties, CP, CG, and CQ for a set of appropriate parameters, and will use this to derive a UC simulator that respects the restrictions of our ledger’s `ExtendPolicy` with the same parameters.

The above policy of the new ledger is similar to the one from [4,5]. What differentiates the new Chronos ledger is an *exportable approximate-time extension*: Every party  $P$  has an associated timestamp  $t$ —which will correspond to this party’s local time in Chronos.  $t$  is adjusted in two ways: (1) in a monotonically increasing manner whenever the local (baseline) clock (of the party) advances, and (2) in a non-monotonic manner (discussed below), at epoch boundaries. (The beginning and the end of epochs is according to the party’s local time  $t$ —an epoch switch occurs whenever the local time increases by  $R_L$  from the (local) timestamp of the previous epoch.) In fact, our ledger records an *extended* local timestamp `timeP` for each party  $P$ , which is the pair `timeP` =  $(e, t)$ , where  $t$  is as above, and  $e$  is the number of non-monotone adjustments to  $t$ , i.e., the number of epoch switches that  $P$  has observed. Note that due to the potential adversarial clock and network influence combined with dynamic participation, we cannot realize a ledger with perfect, monotonically and consistently increasing timestamps. Therefore we weaken the timestamp updates guarantee to allow the adversary to apply a limited shift  $s$ , within some bounding parameter  $\text{shiftLB} \leq s \leq \text{shiftUB}$ , to the non-monotone adjustments that are triggered by epoch-switches. Nevertheless, the ledger enforces that any two alert parties with respective timestamps  $(e, t)$  and  $(e', t')$ , satisfy the constraints  $|t - t'| \leq \text{timeSlack}_{\text{total}}$  and  $|t - t'| \leq \text{timeSlack}_{\text{ep}}$  if  $e = e'$ , and  $|e - e'| \leq 1$  for the respective clock parameters `timeSlackep`, `timeSlacktotal` that define the maximally allowed skewness of parties. Note that this gives the possibility that within an epoch the slack could be potentially different (i.e., much better) than when comparing two timestamps across epochs.

Our main result is stated in the following. We stress that this is a simplified version of the concrete theorem we actually prove, with some of the parameters left implicit. Our proof actually shows that we can obtain security for a much more concrete set of parameters, reflected fully in Theorem 7. Note that such precision concrete-security statements are crucial for the analysis of Nakamoto-style blockchains as one needs to take decisions such as “how many confirmations are sufficient to ensure a given block cannot be inverted?”

**Theorem.** *Our blockchain synchronizer realizes the ledger functionality described above and achieves the clock parameters*

$$\begin{aligned} \text{shiftLB} &= -2\Delta; & \text{shiftUB} &= \Delta; \\ \text{timeSlack}_{\text{total}} &= 2\Delta; & \text{timeSlack}_{\text{ep}} &= \Delta, \end{aligned}$$

where the parameter  $R_L$  is chosen sufficiently large w.r.t. to the security parameter and  $\Delta = \Delta_{\text{net}} + \Delta_{\text{clock}}$ . The parameters of the ledger and of `ExtendPolicy` are instantiated based on the concrete CP, CG, and CQ guarantees of our blockchain synchronizer.

The ledger functionality is fully specified in Section D. A glossary of all parameters of the ledger is provided in Section M.2. The full UC-realization statement with all parameters concretely instantiated (which is the last step in our sequence of proofs) is given in Section L.

**Using Chronos as a Black-box Synchronizer.** Cryptographic protocols can use the abstracted clock and make use of the provided timestamps. When `timeSlackep` = `timeSlacktotal` = 0, and `shiftLB` = `shiftUB` = 0 we obtain an equivalent formulation of the global time/clock of previous works. If shifts are to be expected but slack is still zero, then a protocol must anticipate irregularities at exactly the epoch boundaries. If in addition the skew is not zero, then a higher-level protocol must be further resilient to parties working with



slightly inaccurate time values. Since the irregularities are very predictable and bounded, it is possible to propose generic countermeasures for protocols that need more accurate clocks. For example, a protocol can stall the operation at epoch boundaries and resume at a fixed later time in the new epoch. Additionally, the duration of one logical round of the higher-level protocol can be enlarged to absorb the skew. By the limited shift, the bounded skew, and the guaranteed advancement of the timestamps one can derive concrete guarantees for liveness. Finally, when working in more optimistic (i.e., less adversarial) network models, the time-stamps can be further adjusted to improve accuracy assuming the network parameters (such as expectation and deviation of delay) are known. We give more details in Section E.

### 3.1 Our Protocol

The protocol we present is a Nakamoto-style proof-of-stake based protocol for the so-called semi-synchronous setting; this is the same model used for standard analyses of Bitcoin. In this model, parties have a somewhat accurate common notion of elapsed time (rather than absolute time information) and the network has an upper bound on the delay which is not known to the parties. At a very high-level the protocol attempts to imitate a process which resembles a situation in which state (including time) is continuously passed on to currently alert stakeholders. The honest majority of active stake assumption that is explicit in [14,4] will then ensure that the adversary cannot destroy this state by using his ability to tune participation.

To ease into the main protocol ideas it is useful to imagine a situation in which there is a core of parties with sufficient stake that has been around from the onset of the blockchain. (These parties have a common, albeit somewhat imperfect, understanding of how much time has passed since the protocol started and can contribute this information to the synchronization procedure.) We stress that the continuous or indefinite presence of such parties is not needed in our final protocol which will ensure that the information that these parties would safeguard is passed on to new parties if/when such inaugural parties go to sleep or deregister.

Here is how such an inaugural participant (i.e., a participant who is assigned stake at the outset of the computation by  $\mathcal{F}_{\text{INIT}}$ ) executes the protocol. With access to the provided genesis block, which reveals an initial record  $\mathcal{S}_1 = ((P_1, v_1^{\text{vrf}}, v_1^{\text{kes}}, s_1), \dots, (P_n, v_n^{\text{vrf}}, v_n^{\text{kes}}, s_n))$  that associates each participant  $P_i$ <sup>11</sup> to its chosen public keys used for verification purposes of the staking process and its initial stake  $s_i$ , each party begins the so-called first epoch of the staking procedure and sets its local clock `localTime` to the value 1. The party has to execute a certain set of tasks per round. Note that two inaugural parties have only a somewhat accurate notion of elapsed time and receiving the genesis block might be delayed, it might very well be that a party  $P_1$  has executed three rounds, while  $P_2$  has only executed one so far, or has not even received the genesis block. The bounds on the clock drifts and the network delay however ensure that the difference of the number of completed protocol rounds does not drift too far apart.

A participant’s main task (per round) is to evaluate whether it is elected to produce a block for the current local time, which we refer to as a *slot*. For this, it evaluates a verifiable random function (VRF) on input  $x := \eta_1 \parallel \text{localTime} \parallel \text{TEST}$ , where  $\eta_1$  is a truly random seed provided by  $\mathcal{F}_{\text{INIT}}$ . If the returned value  $y$  is smaller than a threshold value  $T_P^{\text{ep}}$ , which is derived from the stake associated with  $P$ , then the participant is called a slot leader. The threshold is computed to yield a higher probability of slot leadership the higher the stake of the party. The main task of the slot leader is to create a valid block for this slot that contains, as control information (alongside the transactions), the VRF proof of slot leadership, an additional random nonce, and the hash to the head of the chain it connects to. Each block is signed using a key-evolving signature scheme.<sup>12</sup> As typical in these systems, the block is made to extend (essentially) the longest valid chain known to the party. Due to the slightly shifted local clocks, some care has to be taken to not disregard entirely chains that contain blocks in the logical future of a party. However, the chain a party adopts (and computes the ledger state from) at slot `localTime` shall never contain a block with a higher time-stamp.<sup>13</sup>

<sup>11</sup> More precisely,  $P_i$  denotes just a bitstring in the model that formally identifies a machine and is used to identify which keys (and hence stake) are controlled by corrupted machines. Note that we write participant or party instead of machine.

<sup>12</sup> The KES ensures that if a participants gets corrupted, no blocks can be created in retrospect.

<sup>13</sup> Some further care has to be taken in proof of stake to detect chains that try to perform a long-range attack. We describe this in the next section in more detail when we recall the Genesis chain-selection rule.

In addition to the above actions, or if a party is not slot leader, it must play the lottery once more on input  $x' := \eta_1 \parallel \text{localTime} \parallel \text{SYNC}$ . If the party is lucky this time and receives a return value smaller than the threshold (defined shortly), it must emit a so-called synchronization beacon containing the VRF proof and the current time `localTime`. Synchronization beacons are treated similarly to transactions and are contained into blocks if valid. If a party has done all its tasks, it increments `localTime` and waits until the round is over. Except for the generation of synchronization beacons, which is only done in a first fraction of an epoch, the above round procedure iterates over the entire first epoch, where the length of an epoch is  $R$ , a parameter of the protocol. Our security proof shows that this first epoch does result in a blockchain satisfying common prefix, chain growth, and chain quality properties for specific parameters, as long as the leader-election per slot is to the advantage of honest protocol participants.<sup>14</sup>

At the epoch boundary to the second epoch, two important things happen. First the stake-distribution and the epoch randomness change: they are derived from specific blocks contained in the guaranteed common prefix established by the first epoch. In particular, we must ensure that at the time the stake distribution is fixed, the epoch randomness cannot be predicted to ensure the freshness of the slot leader election lottery for the second epoch. The second critical update at the epoch boundary is the local time: each party performs a local-clock adjustment, outlined in Section 5.1, which ensures that after the adjustment parties are still close together, where “close” means within  $\Delta = \Delta_{\text{net}} + \Delta_{\text{clock}}$  (two sources of bounded variance contribute to this: delay and drift) and that performed shifts of the local clock remain small (which is crucial for security). The desired property follows from the common-prefix guarantee (enabling an agreement on beacons), the honest majority assumption (enabling small clock shifts), and the network properties and clock properties (which ensure correlated arrival times). With some additional considerations detailed in Section 5.1, the protocol proceeds executing the above round tasks for the entire second epoch until the next boundary is met. This iterated process, where one epoch bootstraps the next, is backed by an inductive security argument, following previous works [4,14,26], that shows how the overall security is a consequence of the first epoch achieving the desired blockchain properties to serve as a good basis for the second, etc.

The reason to perform a local-clock adjustment is to enable the main goal of our construction: to enable new parties to safely join the system and to determine, just by observing the network and without any further help, an accurate and up-to-date local-clock value and ledger state with respect to the existing honest parties in the system, i.e., being within a  $\Delta$  interval of their clock values and obtaining the same common-prefix, chain-quality and chain-growth guarantees. After this, newly joining parties can start contributing to the security of the system.

The bootstrapping procedure for newcomers is quite involved due to a combination of obstacles: First, the joining party needs to obtain a blockchain that shares some common prefix with the common prefix established by the existing parties. This is achieved by having the joining party listen to the network for some rounds, and picking the “best” chain  $\mathcal{C}$  it sees in the following sense: when compared with any other seen valid chain  $\mathcal{C}'$ ,  $\mathcal{C}$  contains more blocks in an interval of slots of size  $s$  starting from the forking point of  $\mathcal{C}$  and  $\mathcal{C}'$ . We prove that based on the honest-majority assumption, such a densest chain must share a large common prefix with the chains honest parties currently hold. However,  $\mathcal{C}$  could still be adversarially crafted and for example be much longer than what honest parties agreed on by extending into the future, hence a reliable ledger state cannot yet be computed. However, it will become possible once the joining party succeeds in bootstrapping also an accurate time-stamp in the  $\Delta$  interval of honest participants’ timestamps, which is the second obstacle to overcome. After the party is guaranteed to be hooked to a large prefix of the honest parties’ common-prefix, it begins recording all synchronization beacons it receives on the network for a long enough period of time, a parameter of the system. The length of the waiting time is set in order to ensure that, after the newly joining party started listening to the network, the parties at least once seeded the slot-leadership lottery with a fresh nonce that was unpredictable at the time of joining the system. After an additional waiting time, the agreed-upon set of beacons (with proofs referring to the fresh lottery) will be part of the common prefix and eventually be part of what is known to the joining party. We prove that based on this agreement on beacons found in the blockchain, the clock-adjustments procedure by the current

<sup>14</sup> We note that the leader election is per logical slot and honest parties will all pass through the same logical not at the same time, but at related times.

participants in the system can be retraced and will yield a clock adjustment to the newly joining party’s local clock that will directly push it into the interval of existing honest participants’ local clock. At this point, the party runs the normal chain-selection mechanism, essentially cutting off blocks in its logical future and obtains a reliable ledger state as well.

### 3.2 Outline of the Security Arguments

To handle the proof complexity, we divide the proof logically into a sequence of steps. Here we give a very high-level motivation of this modular approach with pointers to the technical sections. Later in Section 6, we give an in-depth introduction to the more concrete security statements.

We first analyze the protocol only for a single epoch in a rather static world, where parties do not disappear. We focus on proving the traditional CP, CG, and CQ guarantees of the protocol for the first epoch. This is needed to establish that the underlying blockchain protocol, Ouroboros Genesis, is able to tolerate small inaccuracies in terms of time-stamps. To be more precise, the only modification that would be needed at this point is that parties buffer future chains and consider them later. Once we can rely on the blockchain properties, we can as a second step analyze the synchronization procedure and prove that no matter what the adversary does, the parties will always stay close together when transitioning from one epoch (i.e., the first) to the next and the logical clock-adjustments are very small. This establishes the base case for a greater inductive argument to establish the full security over the lifetime of the protocol in a static world. The exact sequence of arguments is given in Section 6.4 and the formal theorems and the proofs are given in the technical supplement, Sections F to H.

The first step in moving to the more dynamic world we are aiming at is the analysis of the joining procedure. A party joining the network acts like an observer of the network (i.e., it does not interfere with the protocol) and becomes synchronized after extracting enough information from the network, at which point it can start to be an active protocol participant. This step of the security proof can hence be conducted based on the previous analysis. This is detailed in Section 6.5 and in the technical supplement, Section I provides the formal theorems and proofs.

At this point, we are ready to derive the CP, CG, and CQ guarantees for the entire protocol in a fully dynamic world, where parties join any time, might be temporarily stalled, and disappear unannounced. This can be argued based on a case distinction on different party types (cf. Section 4) and quantify their impact on the security guarantees established above. As it turns out, the most crucial aspect is the joining of new parties, whose security we already understand at this point. This theorem and its proof are given in the technical supplement, Section J.

Moving towards the UC proof, we first observe that the proven chain-growth guarantees (coupled with CP and CQ) do not directly imply liveness of the ledger as in previous works. The reason is that the parties, due to the time-adjustments at the end of an epoch, could (from the viewpoint of an outside observer) move slower than what the observer perceives as the *nominal time* advancement derived from the baseline speed, i.e., a more objective notion of time. For this reason, we prove a concrete relationship between the reported time-stamps by parties and such a reference time. The definition of the reference time and how to connect it to liveness is discussed in Sections 6.3 and 6.7 respectively, the formal statements with proofs are given in technical Section K.

Putting everything together, we can finally instantiate all ledger parameters and describe the simulator. The clue here is that writing down the simulator at this point is conceptually simple: the simulator can in principle emulate the real-world execution perfectly. As long as the ledger parameters are chosen correctly based on the CP, CQ, and CG parameters established above in the fully dynamic world, only a violation of the properties can prevent a correct simulation. This is the topic of the technical Section L.

## 4 Dynamic Participation Model

To support a fine-grained dynamic participation model, we follow the approach of [4] and categorize the parties into *party types*. Recall that the dynamic participation model allows to capture the security of the

Resource	Basic types of <i>honest</i> parties	
	Resource unavailable	Resource available
random oracle $\mathcal{G}_{\text{RO}}$	<i>stalled</i>	<i>operational</i>
network $\mathcal{F}_{\text{N-MC}}$	<i>offline</i>	<i>online</i>
clock $\mathcal{G}_{\text{PERFLCLOCK}}$	<i>time-unaware</i>	<i>time-aware</i>
synchronized state, local time	<i>desynchronized</i>	<i>synchronized</i>
KES capable of signing (w.r.t. local time)	<i>sign-capable</i>	<i>sign-uncapable</i>

**Derived types:**

$$\text{alert} :\Leftrightarrow \text{operational} \wedge \text{online} \wedge \text{time-aware} \wedge \text{synchronized} \wedge \text{sign-capable}$$

$$\text{active} :\Leftrightarrow \text{alert} \vee \text{adversarial} \vee \text{time-unaware}$$

Note: *alert* parties are honest, *active* parties also contain all adversarial parties.

**Fig. 1.** Party types.

protocol in a realistic fashion, by considering that some parties might be stalling their computation, some might accidentally lose network access and hence disappear unannounced, and others might lose track of the passage of time due to some failure. In our model, we formally let the environment be in charge of connecting and disconnecting to its resources: The various basic and derived types of parties are summarized in Figure 1.

For a given point in execution, a party is considered *offline* if it is not registered with the network, otherwise it is considered *online*. A party is *time-aware* if it is registered with the clock, otherwise we call it *time-unaware*. We say that a party is *operational* if it is registered with the random oracle, otherwise we call it *stalled*. Finally, we say that a party is *sign-capable* if the counter in  $\mathcal{F}_{\text{KES}}$  is less or equal to its local time-stamp.

Additionally, an honest party is called *synchronized* if it has been continuously connected to all its resources for a sufficiently long interval to make sure that, roughly speaking, (i) it holds a chain that shares a common prefix with other synchronized parties (synchronized state) and (ii) its local time does not differ by much from other synchronized parties (synchronized time). Our protocol’s resynchronization procedure `JoinProc` will guarantee the party that after executing it for the prescribed number of rounds, it will achieve both properties (i) and (ii) above. In addition, such a party will eventually become sign-capable in future rounds (in case the KES is “evolved” too far into the future due to a de-synchronized time-stamp before joining). We note that an honest party always knows whether it is synchronized or sign-capable and (in contrast to the treatment in [4]), it maintains its synchronization state in a local variable `isSync` and makes its actions depend on it.

Based on these four basic attributes, we define *alert* and *active* parties similarly to [4]. Alert parties are considered the core set of honest parties that have access to all necessary resources, are synchronized and sign-capable. On the other hand, *potentially active* parties (or *active* for short) are those (honest or corrupted) parties that can potentially act (propose a block, send a synchronization beacon) in its current status; in other words, we cannot guarantee their inactivity. Formally, it includes alert parties, corrupted (i.e., adversarial) parties, and moreover any party that is time-unaware (independently of the other attributes; this is because those parties are in particular not capable of evolving their signing keys reliably and hence it cannot be excluded that if they later get corrupted, they might retroactively perform protocol operations in a malicious way).

The definition of a party type is extended from a single point in an execution to a logical slot in a natural way as follows: a party  $P$  is counted as alert (resp. operational, online, time-aware, synchronized, sign-capable) for a slot  $s1$  if the first time its local clock passes through the (logical) slot  $s1$ , it maintains this state *throughout the whole slot*, otherwise it is considered not alert (resp. stalled, offline, time-unaware, desynchronized, sign-uncapable) for  $s1$ . It is considered corrupted (i.e., adversarial) for  $s1$  if it was corrupted by the adversary  $\mathcal{A}$  when its local clock satisfied `localTime`  $\leq$   $s1$ . Finally, it is active for  $s1$  if it is either corrupted for that slot, or it is alert or time-unaware *at any point* during the interval when its local clock for the first time passes through slot  $s1$ .

## 5 Protocol Details

We provide a more in-depth overview of our new protocol (Ouroboros Chronos) in this section and focus on the actions of an alert party (i.e., synchronized and with access to all required resources). Such a party runs the following standard round instructions:

1. Fetch information from the network (procedure `FetchInformation`) over which transactions, beacons, and blocks are sent. The local time is updated via a call to `UpdateTime`. The party locally advances its time-stamp whenever it realizes that a new (local) round has started by a call to `GIMPERFLCLOCK` and comparison to `lastTick`. The code is found in Section C.6 of the technical supplement.
2. Record the arrival times of the synchronization beacons the protocol sends out (call to `ProcessBeacons`). This is discussed in detail in Section 5.1.
3. Process the received chains: as some chains might have been created by parties whose time-stamps are ahead of local time, the future chains are stored in the buffer `futureChains` for later usage. Among the remaining chains, the protocol will decide whether any chain is more preferable than the local chain using a chain-selection rule inspired by Ouroboros Genesis [4] (procedure `SelectChain`) which we thus refer to as the Genesis rule. An important property of the Genesis rule is that chain selection is secure without requiring a moving checkpoint: roughly speaking, a chain  $\mathcal{C}_1$  is preferred over  $\mathcal{C}_2$  if they have a large common history, except possibly the last  $k$  blocks (where  $k$  is some parameter) and  $\mathcal{C}_1$  is longer. If however, they fork even before, chain  $\mathcal{C}_1$  is preferred if its block density is higher compared to  $\mathcal{C}_2$  in a carefully selected interval of size  $s$  slots after the forking point. The details of all the above are described in Sections C.5, C.8 and C.9,
4. Run the main staking procedure (`StakingProcedure`) to evaluate slot leadership, and potentially create and emit a new block or synchronization beacon. Before the main staking procedure is executed, the local state is updated including the current stake distribution (call to `UpdateStakeDist`). The procedures are specified in Sections C.10 and C.6.
5. If the end of the round coincides with the end of an epoch, the synchronization procedure is executed described in Section 5.1.

The code of the entire protocol is given in the technical supplement, Section C. The code of the basic round structure is given in Section C.3 that steers the above tasks. Below we provide details on the most important aspects of the above process.

**Stake distribution and leader election.** A party  $P$  is an eligible slot-leader for a particular slot  $\mathbf{sl}$  in an epoch  $\mathbf{ep}$  if its VRF-output (for an input dependent on  $\mathbf{sl}$ ) is smaller than a threshold value  $T_P^{\mathbf{ep}}$ . The threshold is derived from the (local) stake distribution  $\mathbb{S}_{\mathbf{ep}}$  assigned to an  $\mathbf{ep}$  which in turn is defined by the (local) blockchain  $\mathcal{C}_{\mathbf{loc}}$ , that is we assume an abstract mapping that assigns to a party (identified by an encoding of its public keys) its stake derived as a function of the transactions in  $\mathcal{C}_{\mathbf{loc}}$ , the genesis block, and the epoch the party is currently in. As described above, the stake distribution is only updated once a party enters a new epoch, i.e., once `localTime` mod  $R = 1$ . Say a party enters in epoch  $\mathbf{ep} + 1$ , then the distribution is defined by the state contained in the block sequence up to and including the last block in epoch  $\mathbf{ep} - 1$  (or the genesis block for the first two epochs). Furthermore, the epoch randomness for epoch  $\mathbf{ep} + 1$  (to refresh the lottery) is extracted from the previous randomness and the seeds defined by the first two-thirds of the blocks in epoch  $\mathbf{ep}$  (for the first epoch, the randomness is defined by the genesis block). Both of these updates thus derived based on the (supposedly) established common prefix among participants.

The relative stake of  $P$  in the stake distribution  $\mathbb{S}_{\mathbf{ep}}$  is denoted as  $\alpha_P^{\mathbf{ep}} \in [0, 1]$ . The mapping  $\phi_f(\cdot)$  is defined as

$$\phi_f(\alpha) \triangleq 1 - (1 - f)^\alpha \tag{1}$$

and is parametrized by a quantity  $f \in (0, 1]$  called the *active slots coefficient* [14].

Finally, the threshold  $T_P^{\mathbf{ep}}$  is determined as

$$T_P^{\mathbf{ep}} = 2^{\ell_{\text{VRF}}} \phi_f(\alpha_P^{\mathbf{ep}}), \tag{2}$$

where  $\ell_{\text{VRF}}$  denotes the output length of the VRF (in bits).

Note that by (2), a party with relative stake  $\alpha \in (0, 1]$  becomes a slot leader in a particular slot with probability  $\phi_f(\alpha)$ , independently of all other parties. We clearly have  $\phi_f(1) = f$ , hence  $f$  is the probability that a hypothetical party controlling all 100% of the stake would be elected leader for a particular slot. Furthermore, the function  $\phi$  has an important property called “independent aggregation” [14]:

$$1 - \phi\left(\sum_i \alpha_i\right) = \prod_i (1 - \phi(\alpha_i)). \quad (3)$$

In particular, when leadership is determined according to  $\phi_f$ , the probability of a stakeholder becoming a slot leader in a particular slot is independent of whether this stakeholder acts as a single party in the protocol, or splits its stake among several “virtual” parties.

The technical description of the staking procedure appears in Section C.10. It starts by two calls evaluating the VRF in two different points, using constants `NONCE` and `TEST` to provide domain separation, and receiving  $(y_\rho, \pi_\rho)$  and  $(y, \pi)$ , respectively. The value  $y$  is used to evaluate slot leadership: if  $y < T_p^{\text{ep}}$  then the party is a slot leader and continues by processing its current transaction buffer to form a new block  $B$ . Aside of this application data, each block contains control information. The information includes the proof of leadership  $(y, \pi)$ , additional VRF-output  $(y_\rho, \pi_\rho)$  that influences the epoch-randomness for the next epoch, and the block signature  $\sigma$  produced using  $\mathcal{F}_{\text{KES}}$ . Finally, an updated blockchain  $\mathcal{C}_{\text{loc}}$  containing the new block  $B$  is multicast over the network (note that in practice, the protocol would only diffuse the new block  $B$ ). A slot leader embeds a sequence of valid transactions into a block. As in [4], we abstract block formation and transaction validity into predicates `blockifyOC` and `ValidTxOC`. The function `blockifyOC` takes as input a plain sequence of transactions and outputs a block, whereas `ValidTxOC` takes as input a single transaction and the ledger state. A transaction is said to be valid with respect to the ledger state if and only if it fulfills the predicate. The transaction validity predicate `ValidTxOC` induces a natural transaction validity on blockchain-states that we succinctly denote by the predicate `isvalidstate(st)` that decides that a state is valid if it can be constructed sequentially by adding one transaction at a time and viewing the already added transactions as part of the state.

**Emitting synchronization beacons.** An alert party emits so-called *synchronization beacons* in the first  $R/6$  slots of an epoch `ep`. To be admissible to emit a beacon, the party evaluates the VRF again as for slot-leadership. To obtain an independent evaluation, we use a new constant called `SYNC` to obtain domain separation. If the returned value  $y \leq T_p^{\text{ep, bc}}$ , where in this case we can simply use a linear scaling of the domain, i.e., we define the threshold

$$T_p^{\text{ep, bc}} := 2^{\ell_{\text{VRF}}} \cdot \alpha_p^{\text{ep}}, \quad (4)$$

then the party will create a block header and send it on the broadcast network.<sup>15</sup>

**Embedding synchronization beacons.** Part of the staking procedure is to embed synchronization beacons in the first  $2R/3$  slots of an epoch `ep`. A synchronization beacon is embedded if the creator of the beacon was elected to emit a beacon (according to the current stake distribution in epoch `ep`) in the first  $R/6$  slots of this epoch, and if no other beacon in the chain already specifies the same slot and party identifiers. Like this, an alert party is assured to produce a valid chain according to the validity predicate `IsValidChain` in Section C.5. Note that for a slot leader, we provide for simplicity an extra-predicate `ValidSB` in Section C.5 that allows ensuring that the extension block is valid with respect to beacon inclusion.

**Running the synchronization procedure.** At the end of an epoch, parties run the synchronization procedure based on the beacons recorded in this epoch. We will elaborate on this core procedure of the new protocol in Section 5.1.

<sup>15</sup> Note that there is no need to additionally sign a beacon. Looking ahead, for the synchronization procedure to achieve its goal, we only need agreement on the reported slot numbers (by the respectively elected parties), which is derived from the blockchain, and the guarantees provided by the broadcast functionality. Furthermore, to bound the shift that alert parties experience, it is sufficient that slot numbers reported by alert (and thus synchronized parties) are dominating and are delivered within a reasonable number of rounds after first being emitted.

## 5.1 The Synchronization Procedure

Our main synchronization procedure is based on several logical building blocks. We describe each of them in detail and provide the rationale behind the choices.

- 1.) *Synchronization slots*: Once a party’s local time-stamp reaches a defined synchronization slot for the first time, it will adjust its local time-stamp before moving to the next slot. The protocol will specify the necessary actions for the cases where the local time-stamp is shifted forward or backward. We define the synchronization slots to be the slots with numbers  $i \cdot R$  for  $i \geq 1$  and hence they coincide with the end of an epoch. In a real-world execution (which is a random experiment with discrete steps), we say that a party  $P$  *has passed its synchronization slot*  $i \cdot R$  (e.g., at step  $x$  of the experiment) if it has already concluded its operations in a round where  $P.\text{localTime} = i \cdot R$  holds for the first time. In the code, the synchronization procedure is invoked as the final step in a synchronization slot in Section C.3.
- 2.) *Synchronization Beacons*: In addition to the other messages, the parties in Ouroboros Chronos generate synchronization messages or “beacons” as follows: an alert party  $P$  evaluates the VRF functionality by sending  $(\text{EvalProve}, \text{sid}, \eta_j \parallel P.\text{localTime} \parallel \text{SYNC})$  to  $\mathcal{F}_{\text{VRF}}$  to receive the response  $(\text{Evaluated}, \text{sid}, y, \pi)$ . The beacon message is then defined as the meta-data

$$\text{SB} \triangleq (P.\text{localTime}, P, y, \pi),$$

where  $P.\text{localTime}$  is the current slot number party  $P$  reports and the triple  $(P, y, \pi)$  is the usual attestation of slot leadership by party (or stakeholder)  $P$ . In the code, synchronization beacons are created in the main staking procedure in Section C.10.

- 3.) *Arrival times bookkeeping*: Every party  $P$  maintains an array  $P.\text{Timestamp}_{\text{SB}}(\cdot)$  that assigns to each synchronization beacon  $\text{SB}$  a pair  $(n, \text{flag}) \in \mathbb{N} \times \{\text{final}, \text{temp}\}$ . Assume a beacon  $\text{SB}$  with  $\text{slotnum}(\text{SB}) \in [j \cdot R + 1, \dots, j \cdot R + R/6]$ ,  $j \in \mathbb{N}$  and party  $P'$  is fetched by party  $P$  (for the first time). If the pair  $(\text{slotnum}(\text{SB}), P')$  is new, the recorded arrival time is defined as follows:
  - If  $P$  has already passed synchronization slot  $j \cdot R$  but not yet passed synchronization slot  $(j + 1) \cdot R$ ,  $\text{Timestamp}_{\text{SB}}(\text{SB})$  is defined as the current slot number and the value is considered final, i.e.,  $\text{Timestamp}_{\text{SB}}(\text{SB}) \triangleq (P.\text{localTime}, \text{final})$ .
  - If party  $P$  has not yet passed synchronization slot  $j \cdot R$  (and thus the beacon belongs logically to this party’s next epoch),  $\text{Timestamp}_{\text{SB}}(\text{SB})$  is defined as the current slot number  $P.\text{localTime}$  and the decision is marked as temporary, i.e.,  $\text{Timestamp}_{\text{SB}}(\text{SB}) \triangleq (P.\text{localTime}, \text{temp})$ . This value will be adjusted once this party adjusts its local time-stamp for the next epoch (when arriving at the next synchronization slot  $j \cdot R$ ).

If a party has already received a beacon for the same slot and creator, it will set the arrival time equal to the first one received among those. The process to record arrival times is described in its own algorithm in Section C.7.

- 4.) *The synchronization interval*: the interval based on which the adjustment of the local time-stamp is computed. For a synchronization slot  $i \cdot R$  ( $i \geq 1$ ), its associated synchronization interval is the interval  $I_{\text{sync}}(i) \triangleq [(i - 1) \cdot R + 1, \dots, (i - 1) \cdot R + R/6]$  and hence encompasses the first sixth of the epoch that is now ending.
- 5.) *Emitting Beacons and inclusion into the chain*: An alert party sends out a synchronization beacon during a synchronization interval (i.e., if the current local time reports a slot number that falls into a synchronization interval) if and only if the VRF evaluation  $(\text{EvalProve}, \text{sid}, \eta_j \parallel P.\text{localTime} \parallel \text{SYNC})$  to  $\mathcal{F}_{\text{VRF}}$  returned  $(\text{Evaluated}, \text{sid}, y, \pi)$  with  $y < T_{\text{P}}^{\text{ep}, \text{bc}}$  where  $T_{\text{P}}^{\text{ep}, \text{bc}}$  is the beacon threshold in the current epoch as defined in equation 4. An alert slot leader  $P'$  on the other hand will include any valid synchronization beacon in its new block as long as  $P'.$

The remaining three steps are implemented as part of the core synchronization procedure in Section C.11.

- 6.) *Computing the adjustment evidence:* The adjustment will be computed based on evidence from the set  $\mathcal{S}_i^P$  that is defined with respect to the current view of  $P$  in the execution: Let  $\mathcal{S}_i^P$  contain all beacons  $\mathbf{SB}$  that report a slot number  $\text{slotnum}(\mathbf{SB}) \in [(i-1) \cdot R + 1, \dots, (i-1) \cdot R + R/6]$  (of the synchronization interval) and which are included in a block  $B$  of  $P.C_{\text{loc}}$  that reports a slot number  $\text{slotnum}(B) \leq (i-1) \cdot R + 2R/3$ . Based on these beacons and their recorded arrival times, the shift will be computed. More precisely, if a beacon  $\mathbf{SB}$  is recorded in  $P.C_{\text{loc}}$ , then the arrival time used in the computation will be based on a the valid<sup>16</sup> beacon  $\mathbf{SB}'$  that reports the same slot number and party identity as  $\mathbf{SB}$  and which has arrived first—either as part of some blockchain block or as a standalone message. By our choice of parameters, parties will have assigned an arrival value to any such beacon with overwhelming probability.
- 7.) *Adjusting the local clock:* The shift  $\text{shift}_i^P$  a party  $P$  computes to adjust its clock in synchronization slot  $i \cdot R$  is defined by

$$\text{shift}_i^P \triangleq \text{med} \{ \text{slotnum}(\mathbf{SB}) - \text{Timestamp}(\mathbf{SB}) \mid \mathbf{SB} \in \mathcal{S}_i^P \}.$$

Recall that  $\text{Timestamp}(\mathbf{SB})$  is shorthand for the first element of the pair  $\text{Timestamp}_{\mathbf{SB}}(\mathbf{SB})$ . As we will show, this adjustment ensures that the local time stamps of alert parties report values in a sufficiently narrow interval (depending on the network delay) to provide all protocol properties we need. Furthermore, for each beacon  $\mathbf{SB}$  with  $P.\text{Timestamp}_{\mathbf{SB}}(\mathbf{SB}) = (a, \text{temp})$  and slot number  $\text{slotnum}(\mathbf{SB}) > i \cdot R$  the arrival time is adjusted by  $P.\text{Timestamp}_{\mathbf{SB}}(\mathbf{SB}) \triangleq (a + \text{shift}_i^P, \text{final})$ . This ensures that eventually the arrival times of all beacons that logically belong to epoch  $i + 1$  will be expressed in terms of the newly adjusted local time-stamp computed at synchronization slot  $i \cdot R$ . At this point, the party is further capable of excluding invalid beacons.

- 8.) At the beginning of the next round the party will report a local time equal to  $i \cdot R + \text{shift} + 1$ . If  $\text{shift} \geq 0$ , the party proceeds by emulating its actions for  $\text{shift}$  rounds. If  $\text{shift} < 0$ , the party remains a silent observer (recording arrival times for example) until its local time has advanced to slot  $i \cdot R + 1$  and resumes normally at that round. Note that in this time, an alert party will not revert any previously reported ledger state with overwhelming probability. The reason is that the party will stick to  $C_{\text{loc}}$  during this waiting time and only replace it by longer chains that do not fork by more than  $k$  blocks from  $C_{\text{loc}}$  which is a direct consequence of the security guarantees implied by the Genesis chain-selection rule. (An alert party reverting a previously reported state implies a common-prefix violation.)

## 5.2 The Joining Procedure

**De-Registration and Re-Joining.** If a party is alert, it can lose in several ways its status of being alert. If a party loses access to the random oracle only, then it will still be able to observe the protocol execution and record message arrivals as seen in Section C.1. The main issue is that such a party—when re-joining—will have to retrace what it missed. This is slightly complicated due to the adjustments to the local clock in the course of the execution. However, the party has all reliable information to actually retrace the actions as if it was present as a passive observer all the time. This special procedure `SimulateClockAdjustments` is described in Section C.13. It is invoked as part of procedure `LedgerMaintenance` before performing as an alert party again.

On the other hand, if any alert party loses access to  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  or  $\mathcal{F}_{\text{N-MC}}$  by the respective de-registration queries, or if it joins anew only late in the execution, then it considers itself as de-synchronized. Parties are aware of their synchronization status, and any party that is de-synchronized will have to run through the main joining procedure `JoinProc`.

**Description of JoinProc.** Introducing synchronization slots into the protocol serves the main purpose of enabling a novel joining procedure that newly joining (or resynchronizing) parties can execute to bootstrap an actual reliable time-stamp and ledger state, where a reliable time-stamp is one that lies in the interval of time stamps reported by alert parties. The joining procedure is divided into several phases where the party gathers reliable information, identifies a good synchronization interval and finally applies the shift(s) that will allow it to report a local time-stamp that is sufficiently close to the alert parties in the system. Below

<sup>16</sup> Evaluated using this epoch’s stake distribution.



we give an overview and rationale behind our procedure and formally prove its security in Section 6.5. The code for this procedure is given in Figure 2 containing the procedure `JoinProc` which is invoked as part of `LedgerMaintenance` for newly joining parties. The procedures refer to parameters that are summarized in Table 1 along with their default values.

**Phase A:** A joining party with all resources available invokes the main round procedure triggering the join procedure that first resets the local variables.

**Phase B:** In the second activation upon a `MAINTAIN-LEDGER` command, the party will jump to phase B and continue to do so until and including round  $t_{\text{off}}$ . During this interval, the party applies the Genesis chain selection rule `maxvalid-bg` to filter its incoming chains. It will apply the chain selection rule to all valid chains it receives. Since the party does not have reliable time, it will consider also future chains as valid, as long as they satisfy all remaining validity predicates (cf. Section C.5). As we prove in Lemma 6, at the end of this phase, the party adopts chain  $\mathcal{C}$  that stands in a particularly useful relation to any chain  $\mathcal{C}'$  an alert party adopts. Roughly, the relation says that the point at which the two chains fork is about  $k$  blocks behind the tip of  $\mathcal{C}'$ . This follows from the Genesis chain selection rule and the fact that  $\mathcal{C}'$  is more dense than  $\mathcal{C}$  shortly after the fork. However, this also means that  $\mathsf{P}$  could still hold an extremely long chain served by the adversary (namely, an adversarial extension of an alert party’s chain at some point less than  $k$  blocks behind the tip into the future). On the positive side, the stake distribution used for general validation of blocks and beacons logically associated to the time before the fork are reliable.

**Phase C:** If a party arrives at local time  $t_{\text{off}} + 1$ , it starts with phase C, the gathering phase. The party still filters chains as before, but now processes the arrival times of beacons from the network (or indirectly via the received chains). This phase is parameterized by two quantities: the sum of  $t_{\text{minSync}}$  and  $t_{\text{stable}}$  define the total duration of this round, where intuitively,  $t_{\text{minSync}}$  guarantees that enough arrival times are recorded to compute a reliable estimate of the time-shift, and  $t_{\text{stable}}$  ensures that the blockchain reaches agreement on which (valid) synchronization beacons to use. After this phase, a party can reliably judge valid arrival times.

**Phase D:** The party collects the valid evidence and computes the adjustment based on the first synchronization interval  $I = [(i - 1)R, \dots, (i - 1)R + R/6]$  identified on the blockchain that reports beacons that arrived sufficiently later than the start of phase C (parameter  $t_{\text{pre}}$ ). Party  $\mathsf{P}$  computes the adjustment value that alert parties would do at synchronization slot  $i \cdot R$  based on the recorded beacon arrival times associated with interval  $I$ . The party  $\mathsf{P}$  is done if its adjusted time does not indicate that it should have passed another synchronization slot (and otherwise, the above is repeated with adjusted arrival times of already recorded beacons).

## 6 Analysis Details

### 6.1 Security Assumptions: Alert and Participating Stake Ratio

We begin by setting down notation and defining the conventions we adopt for measuring stake ratios. The following definition is adapted from [4]; the crucial difference is that it refers to the types of parties with respect to a *logical slot* as defined in Section 4.

**Definition 1 (Classes of parties and their relative stake).** *Let  $\mathcal{P}[\mathbf{s1}]$  denote the set of all parties in a logical slot  $\mathbf{s1}$  and let  $\mathcal{P}_{\text{type}}[\mathbf{s1}]$ , for any type of party described in Figure 1 (e.g. alert, active), denote the set of all parties of the respective type in the slot  $\mathbf{s1}$ . For a set of parties  $\mathcal{P}_{\text{type}}[\mathbf{s1}]$ , let  $\mathcal{S}^-(\mathcal{P}_{\text{type}}[\mathbf{s1}]) \in [0, 1]$  (resp.  $\mathcal{S}^+(\mathcal{P}_{\text{type}}[\mathbf{s1}]) \in [0, 1]$ ) denote the minimum (resp., maximum), taken over the views of all alert parties, of the total relative stake of all the parties in  $\mathcal{P}_{\text{type}}[\mathbf{s1}]$  in the stake distribution used for sampling the slot leaders for slot  $\mathbf{s1}$ .*

Looking ahead, we remark that even though we give the general definition above, our protocol will have the desirable property that for all party types and all time slots,  $\mathcal{S}^-(\mathcal{P}_{\text{type}}[\mathbf{s1}]) = \mathcal{S}^+(\mathcal{P}_{\text{type}}[\mathbf{s1}])$  with overwhelming probability, as all the alert parties will agree on the distribution used for sampling slot leaders with overwhelming probability.

**Definition 2 (Alert ratio, participating ratio).** For any logical slot  $\mathbf{s}1$  during the execution, we let: (i.) the alert stake ratio be the fraction  $\mathcal{S}^-(\mathcal{P}_{\text{alert}}[\mathbf{s}1])/\mathcal{S}^+(\mathcal{P}_{\text{active}}[\mathbf{s}1])$ ; and (ii.) the (potentially) participating stake ratio be  $\mathcal{S}^-(\mathcal{P}_{\text{active}}[\mathbf{s}1])$ .

It is instructive to see that the potentially participating stake ratio allows us to infer the ratio of stake belonging to parties that cannot participate in slot  $\mathbf{s}1$ . Intuitively speaking, we will prove the security of our protocol under the assumption that both stake ratios from Definition 2 are sufficiently lower-bounded (the former one by  $1/2 + \varepsilon$ , the latter one by a constant). We remark that it is easy to verify that in particular, such an assumption also implies the existence of alert parties at any point in the execution.

## 6.2 Blockchain Security Properties

We now define the standard security properties of blockchain protocols: *common prefix*, *chain growth* and *chain quality*. These will later be useful as an intermediate step in establishing the UC-security guarantees.

Similarly to [4], we only grant these guarantees to *alert* parties. More importantly for this work, the definitions from [4] need to be adjusted to take into account the fact that the local clocks of the parties are not synchronized. To this end, we choose now to define the properties below with respect to the *logical* timestamps (i.e., slot numbers) contained in blocks, and the local clocks of the parties. Namely, we refer to logical slots below, and a party is considered to *be on the onset* of slot  $\mathbf{s}1$  (or *enter* slot  $\mathbf{s}1$ ) if her local clock just switched to  $\mathbf{s}1$ .

**Common Prefix (CP); with parameters**  $k \in \mathbb{N}$ . The chains  $\mathcal{C}_1, \mathcal{C}_2$  possessed by two alert parties at the onset of the slots  $\mathbf{s}1_1 < \mathbf{s}1_2$  are such that  $\mathcal{C}_1^{[k]} \preceq \mathcal{C}_2$ , where  $\mathcal{C}_1^{[k]}$  denotes the chain obtained by removing the last  $k$  blocks from  $\mathcal{C}_1$ , and  $\preceq$  denotes the prefix relation.

**Chain Growth (CG); with parameters**  $\tau \in (0, 1], s \in \mathbb{N}$ . Consider a chain  $\mathcal{C}$  possessed by an alert party at the onset of a slot  $\mathbf{s}1$ . Let  $\mathbf{s}1_1$  and  $\mathbf{s}1_2$  be two previous slots for which  $\mathbf{s}1_1 + s \leq \mathbf{s}1_2 \leq \mathbf{s}1$ , so  $\mathbf{s}1_2$  is at least  $s$  slots ahead of  $\mathbf{s}1_1$ . Then  $|\mathcal{C}[\mathbf{s}1_1 : \mathbf{s}1_2]| \geq \tau \cdot s$ . We call  $\tau$  the *speed coefficient*.

**Chain Quality (CQ); with parameters**  $\mu \in (0, 1]$  and  $k \in \mathbb{N}$ . Consider any portion of length at least  $k$  of the chain possessed by an alert party at the onset of a slot; the ratio of blocks originating from alert parties is at least  $\mu$ . We call  $\mu$  the chain quality coefficient.

**Existential Chain Quality ( $\exists$ CQ); with parameter**  $s \in \mathbb{N}$ . Consider a chain  $\mathcal{C}$  possessed by an alert party at the onset of a slot  $\mathbf{s}1$ . Let  $\mathbf{s}1_1$  and  $\mathbf{s}1_2$  be two previous slots for which  $\mathbf{s}1_1 + s \leq \mathbf{s}1_2 \leq \mathbf{s}1$ . Then  $\mathcal{C}[\mathbf{s}1_1 : \mathbf{s}1_2]$  contains at least one alertly generated block (i.e., block generated by an alert party).

The first 3 properties are standard, the last one is a slight variant of chain quality fitting better our analysis. For brevity we sometimes write  $\text{CP}(k)$  (resp.,  $\text{CG}(\tau, s)$ ,  $\text{CQ}(\mu, k)$ ,  $\exists\text{CQ}(s)$ ) to refer to these properties.

While these definitions based on the logical time allow us to talk about the logical structure of the forks created by the parties and reuse parts of the technical machinery given in [26,14,4] to analyze it, providing only guarantees based on the logical time would be unsatisfactory, as the parties running Ouroboros Chronos desire persistence and liveness with respect to a more “real-time” notion (that we define in a moment). We will address this translation from logical-time to real-time guarantees later in Section 6.7.

## 6.3 Formal Definitions of Nominal Time and Skew

As discussed earlier, for many of the security arguments, it is very convenient to define a *nominal time* notion:

**Definition 3 (Nominal Time).** Given an execution of Ouroboros Chronos, any prefix of the execution can be mapped deterministically to an integer  $t$ , which we call *nominal time*, as follows: parsing the prefix from genesis and keeping track of the honest party set registered with the imperfect clock functionality (bootstrapped with the set of inaugural alert parties),  $t$  is the number of times the functionality internally switches all flags  $d_{\mathcal{P}}, \mathcal{P} \in \mathcal{P}$  from 1 to 0 until the final step of the execution prefix. (In case no honest party exists in the execution  $t$  is undefined).

Nominal time is a technical definition useful for the analysis. It naturally coincides with the idea of defining a baseline that runs at a certain speed, but where parties have some varying (but bounded) lead ahead of the baseline. For example, if a set of alert parties execute Chronos from the beginning, then nominal time lower bounds the *number* of rounds completed by any of them. Furthermore, by the bounded (absolute) drift enforced by  $\mathcal{G}_{\text{IMPERFLCLOCK}}^{\Delta_{\text{clock}}}$ , the number of locally completed rounds by these alert parties can each be decomposed to be  $t + \delta$  (nominal) rounds, where  $t$  is the baseline, and  $\delta$  is bounded by  $\Delta_{\text{clock}}$ .

We next state a definition that will help us quantify how much parties' (local) timestamps deviate from the nominal time and from each other.

**Definition 4 (Clock skew and  $\text{Skew}_{\Delta}$ ).** *Given an honest party  $P$ , we define its skew in slot  $\mathbf{s1}$  (denoted  $\text{Skew}^P[\mathbf{s1}]$ ) as the difference between  $\mathbf{s1}$  and the nominal time  $t$  when  $P$  enters slot  $\mathbf{s1}$ . For any  $\Delta \geq 0$  and a slot  $\mathbf{s1}$ , we denote by  $\text{Skew}_{\Delta}[\mathbf{s1}]$  the predicate that for all parties that are synchronized in slot  $\mathbf{s1}$ , their skew in this slot differs by at most  $\Delta$ ; formally*

$$\text{Skew}_{\Delta}[\mathbf{s1}] :\Leftrightarrow \left( \forall P_1, P_2 \in \mathcal{P}_{\text{alert}}[\mathbf{s1}] : \left| \text{Skew}^{P_1}[\mathbf{s1}] - \text{Skew}^{P_2}[\mathbf{s1}] \right| \leq \Delta \right) .$$

Note that in the static-registration setting, all honest parties are synchronized (and hence are considered for  $\text{Skew}_{\Delta}[\mathbf{s1}]$ ); the difference will become important in later sections.

## 6.4 Setting with Static Registration

As outlined in Section 3, our first goal is to establish that the properties of (logical-time) common prefix, chain growth, and chain quality are achieved by Ouroboros Chronos when executed in a restricted environment where all parties participate in the protocol run from the beginning and never get deregistered from any of their resources (i.e., from  $\mathcal{G}_{\text{RO}}$ ,  $\mathcal{F}_{\text{N-MC}}$  or  $\mathcal{G}_{\text{IMPERFLCLOCK}}$ ). Similarly to [4], we refer to this setting as the *setting with static registration*; we will drop this assumption later.

### 6.4.1 Single-Epoch Analysis with $\Delta$ -Bounded Skew

Before we can analyze the resynchronization procedure `SyncProc`, we first need to establish the blockchain security properties given in Section 6.2 for Ouroboros Chronos during a single-epoch execution, as the proper functioning of `SyncProc` will inductively depend on these properties being satisfied in the epochs preceding it. Having this inductive structure of the proof in mind, we actually need a security statement for the single-epoch setting in the regime where the predicate  $\text{Skew}_{\Delta}[\mathbf{s1}]$  is satisfied for all slots in that epoch, we refer to this as the setting *with  $\Delta$ -bounded skew*. Looking ahead, this property will be guaranteed in the first epoch thanks to  $\mathcal{F}_{\text{INIT}}$ , and preserved by induction.

The desired properties are established in the following theorem; we give its proof in Appendix F. Recall that  $R$  denotes the epoch length in slots,  $f$  is the active-slot coefficient, let  $\Delta \geq \Delta_{\text{net}} + \Delta_{\text{clock}}$  be the upper bound on the sum of the network delay and clock drift and let  $\tilde{\Delta} \triangleq 2\Delta$ .

**Theorem 1.** *Consider the single-epoch execution of the protocol *Ouroboros-Chronos* in the setting with static registration and  $\Delta$ -bounded skew. Let  $\alpha, \beta \in [0, 1]$  denote a lower bound on the alert ratio and participating ratio throughout this epoch, respectively. If for some  $\epsilon \in (0, 1)$  we have*

$$\alpha \cdot (1 - f)^{\tilde{\Delta}+1} \geq (1 + \epsilon)/2, \tag{5}$$

and the *maxvalid-bg* parameters,  $k$  and  $s$ , satisfy

$$k > 192\tilde{\Delta}/(\epsilon\beta) \quad \text{and} \quad R/6 \geq s = k/(4f) \geq 48\tilde{\Delta}/(\epsilon\beta f) \tag{6}$$

then *Ouroboros-Chronos* achieves the following guarantees:

**Common prefix.** *The probability that it violates the common prefix property with parameter  $k'$  is no more than  $\bar{\epsilon}_{\text{CP}}(k'; R, \Delta, \epsilon) + \bar{\epsilon}_{\text{mv}}$  where*

$$\bar{\epsilon}_{\text{CP}}(k'; R, \Delta, \epsilon) \triangleq \frac{19R}{\epsilon^4} \exp(\tilde{\Delta} - \epsilon^4 k'/18);$$

**Chain growth.** The probability that it violates the chain growth property with parameters  $s' \geq 48\tilde{\Delta}/(\epsilon\beta f)$  and  $\tau_{\text{CG}} = \beta f/16$  is no more than  $\bar{\epsilon}_{\text{CG}}(\tau_{\text{CG}}, s'; R, \epsilon) + \bar{\epsilon}_{\text{mv}}$  where

$$\bar{\epsilon}_{\text{CG}}(\tau_{\text{CG}}, s'; R, \epsilon) \triangleq \frac{s'R^2}{2} \exp(-(\epsilon\beta f)^2 s'/256) ;$$

**Existential chain quality.** The probability that it violates the existential chain quality property with parameter  $s' \geq 12\tilde{\Delta}/(\epsilon\beta f)$  is no more than  $\bar{\epsilon}_{\exists\text{CQ}}(s'; R, \epsilon) + \bar{\epsilon}_{\text{mv}}$  where

$$\bar{\epsilon}_{\exists\text{CQ}}(s'; R, \epsilon) \triangleq (s' + 1)R^2 \exp(-(\epsilon\beta f)^2 s'/64) ;$$

**Chain quality.** The probability that it violates the chain quality property with parameters  $k' \geq 48\tilde{\Delta}/(\epsilon\beta f)$  and  $\mu = \epsilon\beta f/16$  is no more than  $\bar{\epsilon}_{\text{CQ}}(\mu, k'; R, \epsilon) + \bar{\epsilon}_{\text{mv}}$  where

$$\bar{\epsilon}_{\text{CQ}}(\mu, k'; R, \epsilon) \triangleq \frac{kR^2}{2} \exp(-(\epsilon\beta f)^2 k'/256) ;$$

and where  $\bar{\epsilon}_{\text{mv}}$  is a shorthand for the quantity

$$\bar{\epsilon}_{\text{mv}} \triangleq \exp(\ln R - \Omega(k)) + \bar{\epsilon}_{\text{CG}}(\beta f/16, k/(4f)) + \bar{\epsilon}_{\exists\text{CQ}}(k/(4f)) + \bar{\epsilon}_{\text{CP}}(k\beta/64) .$$

In this work, we analyze the concrete security of the protocol. In asymptotic terms, if  $\kappa$  denotes the security parameter, we can treat the chain selection parameters  $k$  and  $s$ , as well as the parameters  $k'$  and  $s'$  of CP, CG, and CQ, as functions in  $\omega(\log k)$  to conclude that the error terms are negligible such that all chain properties hold (for the concrete coefficients  $\tau_{\text{CG}}$  and  $\mu$ ) except with negligible probability. Note that the bound on  $s$  implies that the epoch length  $R$  has a lower bound in  $\omega(\log k)$ , too.

## 6.4.2 Properties of SyncProc

Here we establish two key properties of the resynchronization procedure SyncProc given in Section 5.1 that is being executed by all alert parties on the edge of any two epochs. Consider the following fact, proven in Appendix G for completeness.

**Fact 1** Let  $(a_i)_{i=1}^n$  and  $(b_i)_{i=1}^n$  be two sequences of  $n$  integers each, with the property that  $|a_i - b_i| \leq \Delta$  for all  $i \in [n]$ . Then we also have  $|\text{med}((a_i)_{i=1}^n) - \text{med}((b_i)_{i=1}^n)| \leq \Delta$ .

The above simple statement is at the heart of our analysis of the synchronization procedure SyncProc. Informally, consider an execution of Ouroboros-Chronos over its full lifetime consisting of several epochs, and focus on the edge between epochs  $i$  and  $i + 1$ , where SyncProc is executed. We show that it satisfies two properties:

**SyncProc maintains  $\text{Skew}_\Delta$ .** If (some parametrizations of) CG and CP are not violated up to the end of epoch  $i$ , then  $\text{Skew}_\Delta$  is satisfied in the first slot of epoch  $i + 1$ .

**Bounded shift.** If a lower bound on  $\alpha$ , some parametrization of  $\exists\text{CQ}$ , and  $\text{Skew}_\Delta$  are not violated up to epoch  $i$ , then the shift by which an alert party updates its local clock in SyncProc right before epoch  $i + 1$  satisfies  $|\text{shift}| \leq 2\Delta$ .

We state each of these properties in a formal manner separately as Lemmas 4 and 5 and give their full proofs in Appendix G.

Here we only briefly comment on the proof of the first property, which relies on two intermediate claims: The first is that all alert parties use the same set of synchronization beacons in their execution of the procedure SyncProc between epochs  $\text{ep}$  and  $\text{ep} + 1$ ; the second is that for any fixed beacon  $\text{SB} \in \mathcal{S}_j^{\text{P}_1} = \mathcal{S}_j^{\text{P}_2}$  (in the  $j$ th synchronization slot), the quantity  $\mu(\text{P}_i, \text{SB}) \triangleq \text{Skew}^{\text{P}_i}[\text{s1}] + \text{slotnum}(\text{SB}) - \text{P}_i.\text{Timestamp}(\text{SB})$  will differ by at most  $\Delta$  between any two alert parties  $\text{P}_1$  and  $\text{P}_2$ .

### 6.4.3 Lifting to Multiple Epochs

Theorem 1 gives us security guarantees achieved by Ouroboros Chronos in a single-epoch setting with static stake distribution and perfect randomness. These guarantees can be extended throughout the whole lifetime of the system consisting of many epochs by an inductive argument over epochs, using the properties of  $\mathcal{F}_{\text{INIT}}$  and Theorem 1 for the base case, and the epoch-randomness analysis of [14] together with the properties of SyncProc from Section 6.4.2 (again together with Theorem 1) for the inductive step. We give an informal theorem here and defer the formal statement (Theorem 5) and its proof to Appendix H.

**Theorem 2 (Full-execution security with static registration; informal).** *Consider the setting with static registration. If the assumptions (5) and (6) are satisfied, then the full execution of Ouroboros-Chronos achieves the same guarantees for CP, CG, CQ,  $\exists\text{CQ}$  as given in Theorem 1 except with additional error probability roughly equal to the one induced by the lifting argument in [4].*

While we always target concrete security, an asymptotic summary of this theorem would be that under the same dependencies of the parameters on a security parameter as explained after Theorem 1, the additional error term incurred by the lifting is just negligible.

### 6.5 Newly Joining Parties

In this section we prove that the guarantees on common prefix, chain growth and (existential) chain quality obtained for Ouroboros-Chronos in Section 6.4 remain valid also when new parties join the protocol later during its execution. Again in this section  $\Delta$  denotes the upper bound on the sum of the maximum network delay and maximum clock drift. We use the letter  $t$  to refer to (nominal) time, i.e., to the sequence of execution steps where nominal time is  $t$ .

**Definition 5 (Joining party).** *We say that an honest party  $P$  is joining the protocol execution at time  $t_{\text{join}} > 0$  if  $t_{\text{join}}$  is the nominal time at the point of the execution where  $P$  becomes operational, time-aware and online for the first time.*

Consider an execution of Ouroboros-Chronos over its full lifetime and a joining party  $P_{\text{join}}$  with access to all its resources. Informally speaking, our analysis shows two properties of the joining process of  $P_{\text{join}}$  that hold as long as some parametrizations of CP, CG,  $\exists\text{CQ}$  as well as the assumptions of Theorem 5 remain satisfied throughout the joining process:

1. After Phase B,  $P_{\text{join}}$  will be holding a chain  $C_{\text{join}}$  that satisfies  $C_{\text{alert}}^{\lceil k} \preceq C_{\text{join}}$  with respect to any  $C_{\text{alert}}$  held by an alert party at least  $\Delta$  time steps ago.
2. In Phase D,  $P_{\text{join}}$  correctly identifies an epoch  $i^*$  for which it has collected all the beacons that alert parties had used in their execution of SyncProc after epoch  $i^*$ , and based on these beacons mimics the synchronization procedure so that starting with epoch  $i^* + 1$ ,  $P_{\text{join}}$  does not violate  $\text{Skew}_{\Delta}$  as it becomes alert.

Again, we state each of these properties formally as Lemmas 6 and 7 and give their proofs in Appendix I.

### 6.6 The Dynamic-Availability Setting

Using the above analysis of the joining procedure, we can generalize the results from previous sections to the dynamic availability setting of Section 4, where the parties get arbitrarily registered and deregistered from their resources upon the decision of the environment. We give an informal theorem here and defer the formal statement (Theorem 6) and its proof to Appendix J.

**Theorem 3 (Dynamic availability, informal).** *In the dynamic-availability setting, under the assumptions of Theorem 2 and Lemma 7, Ouroboros-Chronos achieves the same guarantees for CP, CG, CQ and  $\exists\text{CQ}$  as given in Theorem 2 except for a negligible additional error probability that corresponds to violating the assumptions of Lemmas 6 and 7.*

Recall that Lemmas 6 and 7 referred in the above theorem are the formalizations of the argument of Section 6.4.2 and hence their assumptions are informally summarized in that section.

## 6.7 From Logical-Time to Real-Time Guarantees

Recall that eventually, we are interested in a ledger that provides consistency and liveness and they typically follow black-box from the three blockchain properties above. However, since in our protocol, parties emulate a global time themselves, we must make related logical time advancement with the nominal time, which is especially important for liveness. Since parties adjust their timestamps at the boundary of every epoch, an external observer that takes nominal time as the baseline, would conclude that parties are slightly off. To quantify the general relationship, we introduce a concrete discount factor  $\tau_{\text{TG}}$ . The concrete expression is given by Lemma 8 in the technical Section K, and we state the lemma informally here:

**Lemma 1 (Nominal vs. logical time, informal).** *Consider an execution of the full protocol Ouroboros-Chronos in the dynamic-availability setting, let  $P$  be a party that is synchronized between (and including) slots  $s_1$  and  $s_1'$ , let  $t$  and  $t'$  be the nominal times when  $P$  enters slot  $s_1$  and  $s_1'$  for the first time, respectively. Denote by  $\delta s_1$  and  $\delta t$  the respective differences  $|s_1' - s_1|$  and  $|t' - t|$ . Then, under the same assumptions as before, we have  $\delta s_1 \geq \tau_{\text{TG}} \cdot \delta t$  for large enough  $\delta t$ .*

It is important to point out that the  $\tau_{\text{TG}}$  is close to 1 for typical parameter choices and that the lower bound on  $\delta t$  does depend on  $\Delta$  and not on the security parameter. We are ready to state chain-growth with respect to nominal time. Again, the concrete bounds are given in the formal version, Corollary 7 in technical Section K.

**Corollary 1 (Nominal time CG, informal).** *Consider the event that the execution of Ouroboros Chronos under the assumptions as above does not violate property CG with parameters  $\tau \in (0, 1]$ ,  $s \in \mathbb{N}$ . Let  $\tau_{\text{CG, glob}} \triangleq \tau \cdot \tau_{\text{TG}}$ . Consider a chain  $C$  possessed by an alert party at a point in the execution where the party is at an onset of a (local) round and where the nominal time is  $t$ . Let further  $t_1, t_2$ , and  $\delta t$  be such that  $t_1 + \delta t \leq t_2 \leq t$ . Let  $s_{1_1}$  and  $s_{1_2}$  be the last slot numbers that  $P$  reported in the execution when nominal time was  $t_1$  (resp.  $t_2$ ). Then it must hold that  $|C[s_{1_1} : s_{1_2}]| \geq \tau_{\text{CG, glob}} \cdot \delta t$  whenever  $\delta t$  is sufficiently large,*

## 6.8 Composable Guarantees of Ledger and Clock

The final step of our treatment is to show that we get composable security. To this end, we can leverage all the above proven statements and compile all obtained properties into concrete instantiations of ledger parameters. In Section 3 we provided the short form of the full UC statement. The full version is given in Section L where we give the simulation argument and some additional proofs to obtain the stated clock parameters.

## References

1. Handan Kilinc Alper. Ouroboros clepsydra: Ouroboros praos in the universally composable relative time model. Cryptology ePrint Archive, Report 2019/942, 2019. <https://eprint.iacr.org/2019/942>.
2. Hagit Attiya, Amir Herzberg, and Sergio Rajsbaum. Optimal clock synchronization under different delay assumptions (preliminary version). In Jim Anderson and Sam Toueg, editors, *12th ACM PODC*, pages 109–120. ACM, August 1993.
3. Michael Backes, Dennis Hofheinz, Jörn Müller-Quade, and Dominique Unruh. On fairness in simulatability-based cryptographic systems. In Vijay Atluri, Pierangela Samarati, Ralf Küsters, and John C. Mitchell, editors, *Proceedings of the 2005 ACM workshop on Formal methods in security engineering, FMSE 2005, Fairfax, VA, USA, November 11, 2005*, pages 13–22. ACM, 2005.
4. Christian Badertscher, Peter Gazi, Aggelos Kiayias, Alexander Russell, and Vassilis Zikas. Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 18*, pages 913–930. ACM Press, October 2018.
5. Christian Badertscher, Ueli Maurer, Daniel Tschudi, and Vassilis Zikas. Bitcoin as a transaction ledger: A composable treatment. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 324–356. Springer, Heidelberg, August 2017.
6. Iddo Bentov, Rafael Pass, and Elaine Shi. Snow white: Provably secure proofs of stake. Cryptology ePrint Archive, Report 2016/919, 2016. <http://eprint.iacr.org/2016/919>.

7. Gabriel Bracha. An asynchronous  $[(n-1)/3]$ -resilient consensus protocol. In Robert L. Probert, Nancy A. Lynch, and Nicola Santoro, editors, *3rd ACM PODC*, pages 154–162. ACM, August 1984.
8. Jan Camenisch, Robert R. Enderlein, Stephan Krenn, Ralf Küsters, and Daniel Rausch. Universal composition with responsive environments. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 807–840. Springer, Heidelberg, December 2016.
9. Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001.
10. Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. Universally composable security with global setup. In Salil P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 61–85. Springer, Heidelberg, February 2007.
11. Ran Canetti, Kyle Hogan, Aanchal Malhotra, and Mayank Varia. A universally composable treatment of network time. Cryptology ePrint Archive, Report 2017/1256, 2017. <https://eprint.iacr.org/2017/1256>.
12. Jing Chen and Silvio Micali. Algorand. *arXiv preprint arXiv:1607.01341*, 2016.
13. Phil Daian, Rafael Pass, and Elaine Shi. Snow white: Robustly reconfigurable consensus and applications to provably secure proof of stake. In Ian Goldberg and Tyler Moore, editors, *Financial Cryptography and Data Security - 23rd International Conference, FC 2019, Frigate Bay, St. Kitts and Nevis, February 18-22, 2019, Revised Selected Papers*, volume 11598 of *Lecture Notes in Computer Science*, pages 23–41. Springer, 2019.
14. Bernardo David, Peter Gazi, Aggelos Kiayias, and Alexander Russell. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 66–98. Springer, Heidelberg, April / May 2018.
15. Danny Dolev, Joseph Y. Halpern, and H. Raymond Strong. On the possibility and impossibility of achieving clock synchronization. In *16th ACM STOC*, pages 504–511. ACM Press, 1984.
16. Shlomi Dolev and Jennifer L. Welch. Wait-free clock synchronization (extended abstract). In Jim Anderson and Sam Toueg, editors, *12th ACM PODC*, pages 97–108. ACM, August 1993.
17. Shlomi Dolev and Jennifer L. Welch. Self-stabilizing clock synchronization in the presence of byzantine faults (abstract). In James H. Anderson, editor, *14th ACM PODC*, page 256. ACM, August 1995.
18. Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 281–310. Springer, Heidelberg, April 2015.
19. Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol with chains of variable difficulty. Cryptology ePrint Archive, Report 2016/1048, 2016. <http://eprint.iacr.org/2016/1048>.
20. Joseph Y. Halpern, Barbara Simons, H. Raymond Strong, and Danny Dolev. Fault-tolerant clock synchronization. In Robert L. Probert, Nancy A. Lynch, and Nicola Santoro, editors, *3rd ACM PODC*, pages 89–102. ACM, August 1984.
21. Dennis Hofheinz and Joern Mueller-Quade. A synchronous model for multi-party computation and the incompleteness of oblivious transfer. Cryptology ePrint Archive, Report 2004/016, 2004. <http://eprint.iacr.org/2004/016>.
22. Yael Kalai, Yehuda Lindell, and Manoj Prabhakaran. Concurrent composition of secure protocols in the timing model. Cryptology ePrint Archive, Report 2005/036, 2005. <http://eprint.iacr.org/2005/036>.
23. Yael Tauman Kalai, Yehuda Lindell, and Manoj Prabhakaran. Concurrent composition of secure protocols in the timing model. *Journal of Cryptology*, 20(4):431–492, October 2007.
24. Jonathan Katz, Ueli Maurer, Björn Tackmann, and Vassilis Zikas. Universally composable synchronous computation. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 477–498. Springer, Heidelberg, March 2013.
25. T. Kerber, A. Kiayias, M. Kohlweiss, and V. Zikas. Ouroboros crypsinous: Privacy-preserving proof-of-stake. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 984–1001, Los Alamitos, CA, USA, may 2019. IEEE Computer Society.
26. Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 357–388. Springer, Heidelberg, August 2017.
27. Leslie Lamport and P. M. Melliar-Smith. Byzantine clock synchronization. In Robert L. Probert, Nancy A. Lynch, and Nicola Santoro, editors, *3rd ACM PODC*, pages 68–74. ACM, August 1984.
28. Christoph Lenzen, Thomas Locher, and Roger Wattenhofer. Clock synchronization with bounded global and local skew. In *49th FOCS*, pages 509–518. IEEE Computer Society Press, October 2008.
29. Aanchal Malhotra, Matthew Van Gundy, Mayank Varia, Haydn Kennedy, Jonathan Gardner, and Sharon Goldberg. The security of ntp’s datagram protocol. In Aggelos Kiayias, editor, *Financial Cryptography and Data Security - 21st International Conference, FC 2017, Sliema, Malta, April 3-7, 2017, Revised Selected Papers*, volume 10322 of *Lecture Notes in Computer Science*, pages 405–423. Springer, 2017.
30. David L. Mills. *Computer Network Time Synchronization: The Network Time Protocol on Earth and in Space, Second Edition*. CRC Press, 2010.

31. Jesper Buus Nielsen. *On Protocol Security in the Cryptographic Model*. PhD thesis, University of Aarhus, June 2003.
32. Rafail Ostrovsky and Boaz Patt-Shamir. Optimal and efficient clock synchronization under drifting clocks. In Brian A. Coan and Jennifer L. Welch, editors, *18th ACM PODC*, pages 3–12. ACM, May 1999.
33. Rafael Pass, Lior Seeman, and abhi shelat. Analysis of the blockchain protocol in asynchronous networks. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part II*, volume 10211 of *LNCS*, pages 643–673. Springer, Heidelberg, April / May 2017.
34. Rafael Pass and Elaine Shi. Rethinking large-scale consensus. In *30th IEEE Computer Security Foundations Symposium, CSF 2017, Santa Barbara, CA, USA, August 21-25, 2017*, pages 115–129. IEEE Computer Society, 2017.
35. Rafael Pass and Elaine Shi. The sleepy model of consensus. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part II*, volume 10625 of *LNCS*, pages 380–409. Springer, Heidelberg, December 2017.
36. Barbara B. Simons, Jennifer Lundelius Welch, and Nancy A. Lynch. An overview of clock synchronization. In Barbara B. Simons and Alfred Z. Spector, editors, *Fault-Tolerant Distributed Computing*, volume 448 of *LNCS*. Springer, Heidelberg, 1990.
37. T. K. Srikanth and Sam Toueg. Optimal clock synchronization. In Michael A. Malcolm and H. Raymond Strong, editors, *4th ACM PODC*, pages 71–86. ACM, August 1985.

## Appendix and Technical Sections

### A Short Survey of Related Literature

PoS protocols have been proposed as a sustainable and potentially more scalable alternative to the resource-intensive proof-of-work paradigm, but had initially been criticised as a much more limited technology, due to early attacks such as the nothing-at-stake or the stake-bleeding attack. The systematic study of such blockchains, however, has established them as a viable replacement to PoW blockchains, by proving the ability to counter adaptive attacks [14,26], removing the need of checkpointing while assuming dynamic availability [4], adding transaction privacy [25], and, new in our work, removing the dependency on global time for participation-unrestricted environments.

On a different note, as a tool for bootstrapping loose synchronization—local clocks advancing with approximately the same speed—to (approximate) global time, Chronos shares similarities in its goal with *synchronizers*, which occupy an important part of the distributed computing literature. A synchronizer is a sophisticated distributed fault-tolerant protocol which achieves a similar bootstrapping goal, but relies on either knowledge of concrete participation bounds, or at times even knowledge of credentials of the participants. In a nutshell, in typical synchronizer scenarios [15,27,20] it is assumed that the parties have initially (loosely) synchronized clocks and that the clocks advance at about the same speed. It is proved that without setup assumptions, such as a public-key infrastructure that enables digital signatures,  $n$  parties can synchronize their clocks and keep them (loosely) synchronized if and only if no more than  $t < n/3$  of the parties report far-drifting or inconsistent clocks values [15,27,20]. This bound can be improved to  $t < n/2$  by use of existentially unforgeable digital signatures [27]. A number of follow up works have investigated clock synchronization in various settings [17,16,2,37,28,32] and we refer the reader to [36] for a survey (albeit somewhat outdated).

Unfortunately the techniques used in the above traditional synchronizer setting do not translate to the unrestricted participation setting we consider here. Informally, the reason is that these techniques rely on knowledge (and counting) the total number of parties in the system, an assumption which one cannot make in our setting. More concretely, the main tool used by classical synchronizers is counting of messages in combination with signatures to thwart malicious behavior. In the dynamic availability setting, counting messages does not work (or even counting “stake” in the PoS setting), as the parties have no way of knowing how many, or which parties are present at any given time. We note in passing that this is also a major factor that distinguishes the Algorand [12] approach to decentralized consensus from that we follow here. Concretely, Algorand explicitly assumes that parties know (approximately) how many parties are in any committee, and therefore, know how many honest messages they can expect by any such committee. This assumption has been shown to allow Algorand to achieve complete agreement on the whole blockchain after each block; and,



although not explicitly proven we conjecture that techniques from the permissioned synchronizers literature can be employed here to achieve global clock synchronization in the setting of [12]. It is evident, however, that this poses a much stronger assumption than that of dynamic availability leaving open the question of whether PoS-based clock synchronization in the dynamic availability setting is even possible, which we answer affirmatively in our work.

The protocol most closely related to Chronos, which is also the starting point of the underlying blockchain, is Ouroboros Genesis [4]. Notwithstanding, the two solutions have major differences both in terms of their assumptions (access to global time vs. merely similar speed clocks), and in terms of techniques used in the protocol design and in the proof. We refer to Section C for a discussion of these differences. We note that a recent and concurrent work [1] which appeared after our paper (an earlier version of our paper<sup>17</sup> with only minor differences was made public prior to the publication of [1]) also investigated the possibility of using the Ouroboros family of protocols for synchronization of newly joining parties. However [1] lacks a full treatment of dynamic availability and appears to follow a protocol design that we have already argued in the introduction to be insufficient for providing secure synchronization.

## B Model Details

The purpose of this section is to give the details of the underlying (UC) model.

**Relaxed Synchrony.** We use the following version of imperfect local clocks that allow parties to proceed at roughly the same speed, where “roughly” is captured by the upper bound  $\Delta_{\text{clock}}$  on the drift between any two honest parties:

### Functionality $\mathcal{G}_{\text{IMPERFLOCK}}^{\Delta_{\text{clock}}}$

The functionality manages the set  $\mathcal{P}$  of registered identities, i.e., parties  $P = (\text{pid}, \text{sid})$ . It also manages the set  $F$  of functionalities (together with their session identifier). Initially,  $\mathcal{P} := \emptyset$  and  $F := \emptyset$ .

For each identity  $P := (\text{pid}, \text{sid}) \in \mathcal{P}$  it manages bit variables  $d_P$  and  $d_P^{Imp}$ , and a (possibly negative) integer **drift<sub>P</sub>**. For each pair  $(\mathcal{F}, \text{sid}) \in F$  it manages variable  $d_{(\mathcal{F}, \text{sid})} \in \{0, 1\}$  (all these variables are initially set to 0).

*Synchronization:*

- Upon receiving (CLOCK-UPDATE,  $\text{sid}_C$ ) from some party  $P \in \mathcal{P}$  set  $d_P^{Imp} := 1$  and  $d_P := 1$ ; execute *Round-Update* and forward (CLOCK-UPDATE,  $\text{sid}_C, P$ ) to  $\mathcal{A}$ .
- Upon receiving (CLOCK-UPDATE,  $\text{sid}_C$ ) from some functionality  $\mathcal{F}$  in a session  $\text{sid}$  such that  $(\mathcal{F}, \text{sid}) \in F$  set  $d_{(\mathcal{F}, \text{sid})} := 1$ , execute *Round-Update* and return (CLOCK-UPDATE,  $\text{sid}_C, \mathcal{F}$ ) to this instance of  $\mathcal{F}$ .
- Upon receiving (CLOCK-PUSH,  $\text{sid}_C, P$ ) from  $\mathcal{A}$  where party  $P \in \mathcal{P}$ , if  $d_P^{Imp} := 1$  and **drift<sub>P</sub>** <  $\Delta_{\text{clock}}$  then update  $d_P^{Imp} := 0$  and **drift<sub>P</sub>** := **drift<sub>P</sub>** + 1 and return (CLOCK-PUSH-OK,  $\text{sid}_C, P$ ) to  $\mathcal{A}$ . Otherwise ignore the message.
- Upon receiving (CLOCK-GET,  $\text{sid}_C$ ) from any participant  $P$ —including the environment on behalf of a party— or the adversary on behalf of a corrupted party  $P$  (resp. from any ideal—shared or local—functionality  $\mathcal{F}$ ), execute procedure *Round-Update*; return (CLOCK-GET,  $\text{sid}_C, d_P^{Imp}$ ) (resp. (CLOCK-GET,  $\text{sid}_C, d_{(\mathcal{F}, \text{sid})}$ )) to the requestor (where  $\text{sid}$  is the  $\text{sid}$  of the calling instance).

*Procedure Round-Update:* For each session  $\text{sid}$  do: If  $d_{(\mathcal{F}, \text{sid})} := 1$  for all  $\mathcal{F} \in F$  and  $d_P = 1$  for all honest parties  $P = (\cdot, \text{sid}) \in \mathcal{P}$ , then update  $d_{(\mathcal{F}, \text{sid})} := 0$  and update  $d_P := 0$  and **drift<sub>P</sub>** := **drift<sub>P</sub>** – 1 for all parties  $P = (\cdot, \text{sid}) \in \mathcal{P}$ ; for all  $P = (\cdot, \text{sid}) \in \mathcal{P}$  with **drift<sub>P</sub>** < 0 reset  $d_P^{Imp} := 0$  and **drift<sub>P</sub>** := 0.

To improve accessibility, we colored in blue the differences, i.e., the much weaker coordination that the functionality provides compared to the original clock of [24]. In order to appreciate the difference, we sketch the clock from [24] here:

<sup>17</sup> See <https://eprint.iacr.org/2019/838>

### Functionality $\mathcal{G}_{\text{PERFLCLOCK}}$

The functionality manages the set  $\mathcal{P}$  of registered identities, i.e., parties  $P = (\text{pid}, \text{sid})$ . It also manages the set  $F$  of functionalities (together with their session identifier). Initially,  $\mathcal{P} := \emptyset$  and  $F := \emptyset$ .

For each identity  $P := (\text{pid}, \text{sid}) \in \mathcal{P}$  it manages variable  $d_P \in \{0, 1\}$ . For each pair  $(\mathcal{F}, \text{sid}) \in F$  it manages variable  $d_{(\mathcal{F}, \text{sid})} \in \{0, 1\}$  (both bit variables are initially set to 0).

*Synchronization:*

- Upon receiving  $(\text{CLOCK-UPDATE}, \text{sid}_C)$  from some party  $P \in \mathcal{P}$  set  $d_P := 1$ ; execute *Round-Update* and forward  $(\text{CLOCK-UPDATE}, \text{sid}_C, P)$  to  $\mathcal{A}$ .
- Upon receiving  $(\text{CLOCK-UPDATE}, \text{sid}_C)$  from some functionality  $\mathcal{F}$  in a session  $\text{sid}$  such that  $(\mathcal{F}, \text{sid}) \in F$  set  $d_{(\mathcal{F}, \text{sid})} := 1$ , execute *Round-Update* and return  $(\text{CLOCK-UPDATE}, \text{sid}_C, \mathcal{F})$  to this instance of  $\mathcal{F}$ .
- Upon receiving  $(\text{CLOCK-GET}, \text{sid}_C)$  from any participant  $P$ —including the environment on behalf of a party—or the adversary on behalf of a corrupted party  $P$  (resp. from any ideal—shared or local—functionality  $\mathcal{F}$ ), execute procedure *Round-Update*, return  $(\text{CLOCK-GET}, \text{sid}_C, d_P)$  (resp.  $(\text{CLOCK-GET}, \text{sid}_C, d_{(\mathcal{F}, \text{sid})})$ ) to the requestor (where  $\text{sid}$  is the  $\text{sid}$  of the calling instance).

*Procedure Round-Update:* For each session  $\text{sid}$  do: If  $d_{(\mathcal{F}, \text{sid})} := 1$  for all  $\mathcal{F} \in F$  and  $d_P = 1$  for all honest parties  $P = (\cdot, \text{sid}) \in \mathcal{P}$ , then update  $d_{(\mathcal{F}, \text{sid})} := 0$  and  $d_P := 0$  for all parties  $P = (\cdot, \text{sid}) \in \mathcal{P}$ .

**Modeling Peer-to-Peer Communication.** We assume a diffusion network in which all messages sent by honest parties are guaranteed to be fetched by protocol participants after a specific delay  $\Delta_{\text{net}}$ . Additionally, the network guarantees that once a message has been fetched by an honest party, this message is fetched by any other honest party within a delay of at most  $\Delta_{\text{net}}$ , even if the sender of the message is corrupted. This is a slightly stronger guarantee than the multicast-functionality from [5,4] which only guaranteed this bounded delivery for messages sent by honest parties, however it appears to capture well the behavior of peer-to-peer gossiping where honest parties always forward messages they have not already seen. To avoid confusion, we refer to using such a network as broadcasting. We detail the corresponding functionality below:

### Functionality $\mathcal{F}_{\text{N-MC}}^{\Delta_{\text{net}}}$

The functionality is parameterized with a set possible senders and receivers  $\mathcal{P}$ . Any newly registered (resp. deregistered) party is added to (resp. deleted from)  $\mathcal{P}$ .

- **Honest sender multicast.** Upon receiving  $(\text{MULTICAST}, \text{sid}, m)$  from some  $P \in \mathcal{P}$ , where  $\mathcal{P} = \{U_1, \dots, U_n\}$  denotes the current party set, choose  $n$  new unique message-IDs  $\text{mid}_1, \dots, \text{mid}_n$  of the form  $\text{mid}_i = (\text{mid}, i)$ , initialize  $2n$  new variables  $D_{\text{mid}_1} := D_{\text{mid}_1}^{\text{MAX}} \dots := D_{\text{mid}_n} := D_{\text{mid}_n}^{\text{MAX}} := 1$ , a per message-delay  $\Delta_{\text{mid}_i} = \Delta_{\text{net}}$  for  $i = 1, \dots, n$  and set  $\mathbf{M} := \mathbf{M} \parallel (m, \text{mid}_1, D_{\text{mid}_1}, U_1) \parallel \dots \parallel (m, \text{mid}_n, D_{\text{mid}_n}, U_n)$ , and send  $(\text{MULTICAST}, \text{sid}, m, P, (U_1, \text{mid}_1), \dots, (U_n, \text{mid}_n))$  to the adversary.
- **Adversarial sender multicast.** Upon receiving  $(\text{MULTICAST}, \text{sid}, m)$  from some  $P \in \mathcal{P}$  (where  $\mathcal{P} = \{U_1, \dots, U_n\}$  denotes the current party set), do execute it the same way as an honest-sender multicast, with the only difference that  $\Delta_{\text{mid}_i} = \infty$ .
- **Honest party fetching.** Upon receiving  $(\text{FETCH}, \text{sid})$  from  $P \in \mathcal{P}$  (or from  $\mathcal{A}$  on behalf of  $P$  if  $P$  is corrupted):
  1. For all tuples  $(m, \text{mid}, D_{\text{mid}}, P) \in \mathbf{M}$ , set  $D_{\text{mid}} := D_{\text{mid}} - 1$ .
  2. Let  $\mathbf{M}_0^P$  denote the subvector  $\mathbf{M}$  including all tuples of the form  $(m, \text{mid}, D_{\text{mid}}, P)$  with  $D_{\text{mid}} = 0$  (in the same order as they appear in  $\mathbf{M}$ ). Then, delete all entries in  $\mathbf{M}_0^P$  from  $\mathbf{M}$  and in case some  $(m, \text{mid}, D_{\text{mid}}, P)$  is in  $\mathbf{M}_0^P$ , where  $P$  is honest, set  $\Delta_{\text{mid}'} = \Delta_{\text{net}}$  for any  $(m, \text{mid}', D_{\text{mid}'}, P')$  in  $\mathbf{M}$  and replace this record by  $(m, \text{mid}', \min\{D_{\text{mid}'}, \Delta_{\text{net}}\}, P')$ . Finally, send  $\mathbf{M}_0^P$  to  $P$ .
- **Adding adversarial delays.** Upon receiving  $(\text{DELAYS}, \text{sid}, (T_{\text{mid}_{i_1}}, \text{mid}_{i_1}), \dots, (T_{\text{mid}_{i_\ell}}, \text{mid}_{i_\ell}))$  from the adversary do the following for each pair  $(T_{\text{mid}_{i_j}}, \text{mid}_{i_j})$ :

If  $D_{\text{mid}_{i_j}}^{\text{MAX}} + T_{\text{mid}_{i_j}} \leq \Delta_{\text{mid}_{i_j}}$  and  $\text{mid}_{i_j}$  is a message-ID registered in the current  $M$ , set  $D_{\text{mid}_{i_j}} := D_{\text{mid}_{i_j}} + T_{\text{mid}_{i_j}}$  and set  $D_{\text{mid}_{i_j}}^{\text{MAX}} := D_{\text{mid}_{i_j}}^{\text{MAX}} + T_{\text{mid}_{i_j}}$ ; otherwise, ignore this pair.

- **Adversarially reordering messages.** Upon receiving  $(\text{SWAP}, \text{sid}, \text{mid}, \text{mid}')$  from the adversary, if  $\text{mid}$  and  $\text{mid}'$  are message-IDs registered in the current  $M$ , then swap the triples  $(m, \text{mid}, D_{\text{mid}}, \cdot)$  and  $(m, \text{mid}', D_{\text{mid}'}, \cdot)$  in  $M$ . Return  $(\text{SWAP}, \text{sid})$  to the adversary.

**The Genesis Block Distribution with Weak Start Agreement.** In this work, we not only allow that parties’ local time-stamps might shift apart over the course of an execution, we also model that the genesis block is received by initial stakeholders in a delayed fashion (such that one party already starts staking, but another one is still waiting for the genesis block). To this aim, we weaken the functionality  $\mathcal{F}_{\text{INIT}}$  described in [4] to allow for bounded offsets in starting times. Namely,  $\mathcal{F}_{\text{INIT}}^{\Delta_{\text{net}}}$  does not enforce that all honest stakeholders receive the created genesis block in the same round, but merely guarantees delivery differences of  $\Delta_{\text{net}}$ , just as the message delays above (the offsets are under adversarial control). In fact, we find it reasonable to think of the genesis block delay as being considered similar to a “message delay”, i.e., we model that the genesis blocks “arrives with a delay”. The details of the  $\mathcal{F}_{\text{INIT}}^{\Delta_{\text{net}}}$  functionality appear below:

More concretely, we allow the adversary to define the offsets upon the first activation to the functionality. For the sake of convenience, we consider this initial offset query to the adversary as restricting (and prefix the query with the keyword **Respond**) as defined by Camenisch et al. [8] which means that the adversary is required to answer this query immediately (and hence the offsets can technically be seen as chosen “when the  $\mathcal{F}_{\text{INIT}}$  is created”).<sup>18</sup> The functionality makes sure that at the onset of the execution, the genesis block is actually created with the keying material from the initial stakeholders. As long as this creation is not yet complete, the baseline cannot advance (recall that ideal functionalities are allowed to proceed at the baseline speed). Once the creation is complete, which defines the “big bang” for this execution, the clock is released, and at this point, the offsets steer when a initial stakeholder receives genesis block and start working.

#### Functionality $\mathcal{F}_{\text{INIT}}^{\Delta_{\text{net}}}$

The functionality  $\mathcal{F}_{\text{INIT}}$  is parameterized by the delay  $\Delta_{\text{net}}$ , the set  $P_1, \dots, P_n$  of initial stakeholders  $n$  and their respective stakes  $s_1, \dots, s_n$ . It additionally stores  $n$  variables  $\text{offset}_i$ , one for each stakeholder  $P_i$  to steer when they receive the genesis block (once ready), a variable  $\text{ready}$  to steer when the genesis block is fully generated. It maintains the set of registered parties  $\mathcal{P}$ .

- On the first activation of the functionality, send  $(\text{Respond}, (\text{DefineOffset}, \text{sid}))$  to the adversary. Upon receiving the response  $(\text{DefineOffset}, \text{sid}, o_1, \dots, o_n)$  where  $o_1 \in [0, \dots, \Delta_{\text{net}}]$ , set  $\text{offset}_i := o_i$ . It further sets  $\text{ready} \leftarrow 0$  and registers with the clock. Proceed handling the first input as specified in the following.
- On receiving any input by a registered party, do the following case distinction:
  - If  $\text{ready} = 0$  and the message is a request from some initial stakeholder  $P = P_i$ ,  $i \in [n]$ , of the form  $(\text{ver\_keys}, \text{sid}, P, v^{\text{vrf}}, v^{\text{kes}})$ , then  $\mathcal{F}_{\text{INIT}}$  stores the verification keys tuple  $(P_i, v_i^{\text{vrf}}, v_i^{\text{kes}})$  and acknowledges its receipt. If some of the registered public keys are equal, it outputs an error and halts. Otherwise, it samples and stores a random value  $\eta_1 \xleftarrow{\$} \{0, 1\}^\lambda$  and constructs a genesis block  $(S_1, \eta_1)$ , where  $S_1 = ((P_1, v_1^{\text{vrf}}, v_1^{\text{kes}}, s_1), \dots, (P_n, v_n^{\text{vrf}}, v_n^{\text{kes}}, s_n))$ . If all initial stakeholders have made their requests, then  $\mathcal{F}_{\text{INIT}}$  de-registers from the clock to release the execution.
  - If  $\text{ready} = 1$ , then do the following
    - \* If the currently received input is the  $k$ th request of the form  $(\text{genblock\_req}, \text{sid}, P)$  from an initial stakeholder  $P = P_i$  (for some  $i \in [n]$ ), if  $k > \text{offset}_i$ ,  $\mathcal{F}_{\text{INIT}}$  sends  $(\text{genblock}, \text{sid}, (S_1, \eta_1))$  to  $P$ .

<sup>18</sup> In case the query would be not be restricting, this would incur a slight but rather artificial complication of the initialization procedure of the protocol as we would have to take into account that the very first activated honest protocol participant (and only this one) will actually lose its activation token right away. Defining this query to be restricting is not crucial for our treatment.

\* If it is a request from a party that is not initial stakeholder then return ((genblock, sid, (S<sub>1</sub>, η<sub>1</sub>)), Running) to the requestor. // The case when the setup is completed and potentially, the protocol has advanced a lot and the party must bootstrap from this genesis block.

## B.1 Additional Functionalities Used in the Proof

The protocol makes use of a VRF (verifiable random function) functionality  $\mathcal{F}_{\text{VRF}}$ , a KES (key-evolving signature) functionality  $\mathcal{F}_{\text{KES}}$ , a (global) random oracle functionality  $\mathcal{G}_{\text{RO}}$ . We use the random oracle as the idealization of a hash function. We use the strongest form of a global random oracle to express that our new consensus algorithm does not need any kind of programmability or query restrictions (and the result using a local random oracle is implied). The idealizations  $\mathcal{F}_{\text{VRF}}$  and  $\mathcal{F}_{\text{KES}}$  are shown to be realizable under standard assumptions or an additional random oracle in [14].

Recall that in order to reflect the synchronization and liveness pattern properly in UC, ideal signing or verification (in both, KES or VRF) is a local operation performed by a slot-leader and hence one invokes the proposed formalism of Camenisch et al. [8] and declare signing request as *restricting*, which means that although activated to provide a signature, the adversary has to provide the answer, i.e. the signature string for example, immediately (no other output to another protocol machine is allowed) and return the activation token back to the functionality. Such responsive queries are prefixed with the keyword Respond in outputs to the adversary.

**Verifiable Random Functions.** The idealized VRF functionality is defined below:

**Functionality  $\mathcal{F}_{\text{VRF}}$** 

$\mathcal{F}_{\text{VRF}}$  interacts with its set of registered parties  $\mathcal{P}$  (denoted by  $U_1, \dots, U_{|\mathcal{P}|}$ ) as follows:

- **Key Generation.** Upon receiving a message (KeyGen, sid) from a stakeholder  $U_i$ , hand (Respond, KeyGen, sid,  $U_i$ ) to the adversary. Upon receiving (VerificationKey, sid,  $U_i$ ,  $v$ ) from the adversary, if  $U_i$  is honest, verify that  $v$  is unique, record the pair  $(U_i, v)$  and return (VerificationKey, sid,  $v$ ) to  $U_i$ . Initialize the table  $T(v, \cdot)$  to empty.
- **Malicious Key Generation.** Upon receiving a message (KeyGen, sid,  $v$ ) from  $\mathcal{S}$ , verify that  $v$  has not being recorded before; in this case initialize table  $T(v, \cdot)$  to empty and record the pair  $(\mathcal{S}, v)$ .
- **VRF Evaluation.** Upon receiving a message (Eval, sid,  $m$ ) from  $U_i$ , verify that some pair  $(U_i, v)$  is recorded. If not, then ignore the request. Then, if the value  $T(v, m)$  is undefined, pick a random value  $y$  from  $\{0, 1\}^{\ell_{\text{VRF}}}$  and set  $T(v, m) = (y, \emptyset)$ . Then output (Evaluated, sid,  $y$ ) to  $U_i$ , where  $y$  is such that  $T(v, m) = (y, S)$  for some  $S$ .
- **VRF Evaluation and Proof.** Upon receiving a message (EvalProve, sid,  $m$ ) from  $U_i$ , verify that some pair  $(U_i, v)$  is recorded. If not, then ignore the request. Else, send (Respond, EvalProve, sid,  $U_i$ ,  $m$ ) to the adversary. Upon receiving (EvalProve, sid,  $m$ ,  $\pi$ ) from the adversary, if value  $T(v, m)$  is undefined, verify that  $\pi$  is unique, pick a random value  $y$  from  $\{0, 1\}^{\ell_{\text{VRF}}}$  and set  $T(v, m) = (y, \{\pi\})$ . Else, if  $T(v, m) = (y, S)$ , set  $T(v, m) = (y, S \cup \{\pi\})$ . In any case, output (Evaluated, sid,  $y$ ,  $\pi$ ) to  $U_i$ .
- **Malicious VRF Evaluation.** Upon receiving a message (Eval, sid,  $v$ ,  $m$ ,  $\pi$ ) from  $\mathcal{S}$  for some  $v$ , do the following. First, if  $(\mathcal{S}, v)$  is recorded and  $T(v, m)$  is undefined, then choose a random value  $y$  from  $\{0, 1\}^{\ell_{\text{VRF}}}$  and set  $T(v, m) = (y, S)$  and output (Evaluated, sid,  $y$ ) to  $\mathcal{S}$ . The same is performed in case  $(U_i, v)$  is recorded and  $U_i$  corrupted. Else, if  $T(v, m) = (y, S')$  for some  $S' \neq \emptyset$ , union  $S$  to  $S'$  and output (Evaluated, sid,  $y$ ) to  $\mathcal{S}$ , else ignore the request.
- **Verification.** Upon receiving a message (Verify, sid,  $m$ ,  $y$ ,  $\pi$ ,  $v'$ ) from some party  $P$ , send (Respond, Verify, sid,  $m$ ,  $y$ ,  $\pi$ ,  $v'$ ) to the adversary. Upon receiving (Verified, sid,  $m$ ,  $y$ ,  $\pi$ ,  $v'$ ) from the adversary do:
  1. If  $v' = v$  for some  $(\cdot, v)$  and the entry  $T(v, m)$  equals  $(y, S)$  with  $\pi \in S$ , then set  $f = 1$ .
  2. Else, if  $v' = v$  for some recorded pair of the form  $(\cdot, v)$ , but no entry  $T(v, m)$  of the form  $(y, \{\dots, \pi, \dots\})$  is recorded, then set  $f = 0$ .

3. Else, initialize the table  $T(v', \cdot)$  to empty, and set  $f = 0$ .  
Output  $(\text{Verified}, sid, m, y, \pi, f)$  to  $P$ .

**Key-Evolving Signatures.** The idealized KES module is defined below:

**Functionality  $\mathcal{F}_{\text{KES}}$**

$\mathcal{F}_{\text{KES}}$  is parameterized by the total number of signature updates  $T$ , interacting with a signer  $U_S$  and registered parties in  $\mathcal{P}$  (denoted by  $U_1, \dots, U_{|\mathcal{P}|}$ ) as follows:

- **Key Generation.** Upon receiving a message  $(\text{KeyGen}, sid, U_S)$  from a stakeholder  $U_S$ , send  $(\text{KeyGen}, sid, U_S)$  to the adversary. Upon receiving  $(\text{VerificationKey}, sid, U_S, v)$  from the adversary, verify that  $v$  is unique and send  $(\text{VerificationKey}, sid, v)$  to  $U_S$ , record the triple  $(sid, U_S, v)$  and set counter  $k_{\text{ctr}} = 1$ .
  - **Sign and Update.** Upon receiving a message  $(\text{USign}, sid, U_S, m, j)$  from  $U_S$ , verify that  $(sid, U_S, v)$  is recorded for some  $sid$  and that  $k_{\text{ctr}} \leq j \leq T$ . If not, then ignore the request. Else, set  $k_{\text{ctr}} = j + 1$  and send  $(\text{Respond}, (\text{Sign}, sid, U_S, m, j))$  to the adversary. Upon receiving  $(\text{Signature}, sid, U_S, m, j, \sigma)$  from the adversary, verify that no entry  $(m, j, \sigma, v, 0)$  is recorded. If it is, then output an error message to  $U_S$  and halt. Else, send  $(\text{Signature}, sid, m, j, \sigma)$  to  $U_S$ , and record the entry  $(m, j, \sigma, v, 1)$ .
  - **Signature Verification.** Upon receiving a message  $(\text{Verify}, sid, m, j, \sigma, v')$  from some stakeholder  $U_i$  do:
    1. If  $v' = v$  and the entry  $(m, j, \sigma, v, 1)$  is recorded, then set  $f = 1$ . (This condition guarantees completeness: If the verification key  $v'$  is the registered one and  $\sigma$  is a legitimately generated signature for  $m$ , then the verification succeeds.)
    2. Else, if  $v' = v$ , the signer is not corrupted, and no entry  $(m, j, \sigma', v, 1)$  for any  $\sigma'$  is recorded, then set  $f = 0$  and record the entry  $(m, j, \sigma, v, 0)$ . (This condition guarantees unforgeability: If  $v'$  is the registered one, the signer is not corrupted, and never signed  $m$ , then the verification fails.)
    3. Else, if there is an entry  $(m, j, \sigma, v', f')$  recorded, then let  $f = f'$ . (This condition guarantees consistency: All verification requests with identical parameters will result in the same answer.)
    4. Else, if  $j < k_{\text{ctr}}$ , let  $f = 0$  and record the entry  $(m, j, \sigma, v, 0)$ . Otherwise, if  $j = k_{\text{ctr}}$ , hand  $(\text{Verify}, sid, m, j, \sigma, v')$  to the adversary. Upon receiving  $(\text{Verified}, sid, m, j, \phi)$  from the adversary let  $f = \phi$  and record the entry  $(m, j, \sigma, v', \phi)$ . (This condition guarantees that the adversary is only able to forge signatures under keys belonging to corrupted parties for time periods corresponding to the current or future slots.)
- Output  $(\text{Verified}, sid, m, j, f)$  to  $U_i$ .

## C A Complete Description of the Protocol

The purpose of this section is to specify more formally the code of the Chronos protocol for more clarity with respect to the security claims. We present the ledger protocol following the structure of other ledger protocols in UC [5,4].

### C.1 The Main Protocol Instance

Ouroboros Chronos is a ledger-protocol and the main protocol dispatches to the relevant sub-processes. It accepts three kinds of input: inputs in order to register the party to the required setup and which models the dynamic availability of parties.

**Registration.** A party  $P$  can only start operation once it is registered to all resources. The registration handling is given in Section C.2. The protocol will initialize a party  $P$ 's local time  $P.\text{localTime}$  to 0 and the protocol is aware that it is not synchronized (since existing participants might be far off) and sets  $P.\text{isSync}$  to false. Finally, in order to be able to recognize a new round, the party maintains a variable  $\text{lastTick}$  that stores the most recent tick from  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  (either 0 or 1).

Second, the ledger-specific queries to submit new transactions (SUBMIT), to read the ledger state (READ), to read the exported time value, and finally to make parties work and proceed with their “real round operations” to advance and maintain the ledger.

**Submit transactions.** As in [4] parties take as inputs transactions that serve as the inputs to the ledger.

**Read state.** As in [4], the ledger protocol exports a stable ledger state to the environment (implemented as a certain prefix of the longest chain of a party).

**Read time.** A novelty compared to Ouroboros Genesis, where the global clock showed the global time to all parties, Ouroboros Chronos will export a feature to read the logical protocol time. The exact guarantees on this “new clock”, in particular the skew between reported times and the offset to a related “real-time” notion are given first as properties and then by the ideal ledger-functionality that Ouroboros Chronos realizes.

**Activations requests.** Note that MAINTAIN-LEDGER inputs are the activations that “make the parties work” and perform their round actions in the main procedure `LedgerMaintenance` later.

Finally, the protocol allows a caller controlled access to the features of the global shared setups through this protocol instance and to directly steer the registration status (to model the state in which a party is). We present the relevant sub-protocols and procedures in handling all these calls in the sequel. A summary of the state variables appears in Section M.1. Additionally, the protocol is designed to allow for a predictable number of activations per round and party to allow for a guaranteed advancement (independent of the adversarial actions) as defined in [4].

**Technical remark: handling interrupts in a UC protocol.** As a general paradigm to write the ledger protocol as a UC protocol, we follow the approach taken in [4] to simplify the treatments with interrupts in UC. Note that a protocol command might consists of a sequence of operations. In UC, certain operations, such as sending a message to another party or just the inability to conclude a task because a resource is unavailable, result into the protocol machine having to lose its activation. Thus, one needs a mechanism for ensuring that a party that loses the activation in the middle of such a multi-step command is able to resume and complete this command.

The general mechanism is to introduce an anchor  $a$  that stores a pointer to the current operation; the protocol associates each anchor with an input  $I$ , so that when such an input is received (again) it directly jumps to the stored anchor, executes the next operation(s) and updates (increases) the anchor before releasing the activation. We refer to execution in such a manner as *I-interruptible*.

**Protocol** `Ouroboros-Chronosk`( $P, \text{sid}; \mathcal{G}_{\text{LEDGER}}, \mathcal{G}_{\text{IMPERFLCLOCK}}, \mathcal{G}_{\text{RO}}, \mathcal{F}_{\text{N-MC}}^{\Delta}$ )

**Global Variables:**

- Read-only (parameters):  $R, k, f, s, t_{\text{off}}, t_{\text{stable}}, t_{\text{minSync}}, t_{\text{pre}}$
- Read-write:  $v_{\text{p}}^{\text{vrf}}, v_{\text{p}}^{\text{kes}}, \text{localTime}, \text{ep}, \text{s1}, \mathcal{C}_{\text{loc}}, T_{\text{p}}^{\text{ep}}, \text{isInit}, t_{\text{work}}, \text{buffer}, \text{futureChains}, \text{lastTick}, \text{isSync}, \text{EpochUpdate}(\cdot), \text{fetchCompleted}, \text{lastTimeAlert}, \text{Timestamp}_{\text{SB}}(\cdot)$ . (recall that we use  $\text{Timestamp}(\cdot)$  to denote the first (and numerical) element of the pair  $\text{Timestamp}_{\text{SB}}(\cdot)$ )

**Registration/Deregistration:**

- Upon receiving input `(REGISTER,  $\mathcal{R}$ )`, where  $\mathcal{R} \in \{\mathcal{G}_{\text{LEDGER}}, \mathcal{G}_{\text{IMPERFLCLOCK}}, \mathcal{G}_{\text{RO}}\}$  execute protocol `Registration-Chronos( $P, \text{sid}, \text{Reg}, \mathcal{R}$ )`.
- Upon receiving input `(DE-REGISTER,  $\mathcal{R}$ )`, where  $\mathcal{R} \in \{\mathcal{G}_{\text{LEDGER}}, \mathcal{G}_{\text{IMPERFLCLOCK}}, \mathcal{G}_{\text{RO}}\}$  execute protocol `Deregistration-Chronos( $P, \text{sid}, \text{Reg}, \mathcal{R}$ )`.
- Upon receiving input `(IS-REGISTERED,  $\text{sid}$ )` return `(REGISTER,  $\text{sid}$ , 1)` if the local registry `Reg` indicates that this party has successfully completed a registration with  $\mathcal{R} = \mathcal{G}_{\text{LEDGER}}$  (and did not de-register since then). Otherwise, return `(REGISTER,  $\text{sid}$ , 0)`.

**Interacting with the Ledger:**

Upon receiving a ledger-specific input  $I \in \{(\text{SUBMIT}, \dots), (\text{READ}, \dots), (\text{MAINTAIN-LEDGER}, \dots)\}$  verify first that all resources are available. **If** not all resources are available, **then** ignore the input; **else** (i.e., the party is operational and time-aware) execute one of the following steps depending on the input  $I$ :

- **If**  $I = (\text{SUBMIT}, \text{sid}, \text{tx})$  **then** set  $\text{buffer} \leftarrow \text{buffer} \parallel \text{tx}$ , and send  $(\text{MULTICAST}, \text{sid}, \text{tx})$  to  $\mathcal{F}_{\text{N-MC}}^\Delta$ .
- **If**  $I = (\text{MAINTAIN-LEDGER}, \text{sid}, \text{minerID})$  **then** invoke protocol  $\text{LedgerMaintenance}(\mathcal{C}_{\text{loc}}, \text{P}, \text{sid}, k, s, R, f)$ ; if  $\text{LedgerMaintenance}$  halts **then** halt the protocol execution (all future input is ignored).
- **If**  $I = (\text{READ}, \text{sid})$  **then** invoke protocol  $\text{ReadState}(k, \mathcal{C}_{\text{loc}}, \text{P}, \text{sid}, R, f)$ .
- **If**  $I = (\text{EXPORT-TIME}, \text{sid})$  **then** do the following: if  $\text{isSync}$  or  $\text{isInit}$  is false, then return  $(\text{EXPORT-TIME}, \text{sid}, \perp)$  to the caller. Otherwise call  $\text{UpdateTime}(\text{P}, R)$  and do:
  1. Define  $e$  to be the highest value s.t.  $\text{EpochUpdate}(e) = \text{Done}$ .
  2. Return  $(\text{EXPORT-TIME}, \text{sid}, (e, \text{localTime}))$  to the caller.

#### Handling calls to the shared setup:

- Upon receiving  $(\text{CLOCK-GET}, \text{sid}_C)$  forward it to  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  and output  $\mathcal{G}_{\text{IMPERFLCLOCK}}$ 's response.
- Upon receiving  $(\text{CLOCK-UPDATE}, \text{sid}_C)$ , record that a clock-update was received in the current round. If the party is registered to all its setups, then do nothing further. Otherwise, do the following operations *before concluding this round*:
  1. If this instance is currently time-aware but otherwise stalled or offline, then call  $\text{UpdateTime}(\text{P}, R)$  to update  $\text{localTime}$  and evolve the KES signing key by sending  $(\text{USign}, \text{sid}, \text{P}, 0, \text{localTime})$  to  $\mathcal{F}_{\text{KES}}$ . If the party has passed a synchronization slot, then set  $\text{isSync} \leftarrow \text{false}$ .
  2. If this instance is only stalled but  $\text{isSync} = \text{true}$ , then additionally execute  $\text{FetchInformation}(\text{P}, \text{sid})$ , extract all new synchronization beacons  $\mathcal{B}$  from the fetched chains and record their arrival times and set  $\text{fetchCompleted} \leftarrow \text{true}$ . Also, any unfinished interruptible execution of this round is marked as completed.
  3. Forward  $(\text{CLOCK-UPDATE}, \text{sid}_C)$  to  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  to finally conclude the round.
- Upon receiving  $(\text{EVAL}, \text{sid}_{\text{RO}}, x)$  forward the query to  $\mathcal{G}_{\text{RO}}$  and output  $\mathcal{G}_{\text{RO}}$ 's response.

## C.2 Registration, De-registration

A party  $\text{P}$  needs access to all its resources in order to start operation. Once it is registered to all resources it is able to perform basic operations. In contrast to Ouroboros Genesis, the protocol will initialize a party  $\text{P}$ 's local time  $\text{P.localTime}$  to 0. Furthermore, the protocol is aware that it is not synchronized (since existing participants might be far off) and sets  $\text{P.isSync}$  to false. Finally, in order to be able to recognize a new round, the party maintains a variable  $\text{lastTick}$  that stores the most recent tick from  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  (either 0 or 1).

#### Protocol Registration-Chronos( $\text{P}, \text{sid}, \text{Reg}, \mathcal{G}$ )

- 1: **if**  $\mathcal{G} \in \{\mathcal{G}_{\text{IMPERFLCLOCK}}, \mathcal{G}_{\text{RO}}\}$  **then** send  $(\text{REGISTER}, \text{sid})$  to  $\mathcal{G}$ , set registration status to registered with  $\mathcal{G}$ , and output the valued received by  $\mathcal{G}$ .
- 2: **end if**
- 3: **if**  $\mathcal{G} = \mathcal{G}_{\text{LEDGER}}$  **then**
- 4:   **if** the party is not registered with  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  or  $\mathcal{G}_{\text{RO}}$  **then** or already registered with all setups ignore this input
- 5:   **else**
- 6:     **for** each  $\mathcal{F} \in \{\mathcal{F}_{\text{INIT}}^\Delta, \mathcal{F}_{\text{VRF}}, \mathcal{F}_{\text{KES}}\}$  **do**
- 7:       Send  $(\text{REGISTER}, \text{sid})$  to  $\mathcal{F}$ , set its registration status as registered with  $\mathcal{F}$ , but do not output the received values.
- 8:     **end for**
- 8:   Send  $(\text{CLOCK-GET}, \text{sid}_C)$  to  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  and receive  $(\text{CLOCK-GET}, \text{sid}_C, \text{tick})$  and set  $\text{lastTick} \leftarrow \text{tick}$
- 9:   Send  $(\text{REGISTER}, \text{sid})$  to  $\mathcal{F}_{\text{N-MC}}^\Delta$ .
- 10:   Set  $\text{localTime} := 0$  and  $\text{isSync} \leftarrow \text{false}$ .

```

11:     If this is the first registration invocation for this ITI, then set isInit  $\leftarrow$  false.
12:     Output (REGISTER, sid, P) once completing the registration with all the above resources  $\mathcal{F}$ .
    end if
end if

```

De-registration is then the analogous action of setting variables to the initial values and record the synchronization status correctly:

**Protocol Deregistration-Chronos(P, sid, Reg,  $\mathcal{G}$ )**

```

1: If the party is alert, set lastTimeAlert  $\leftarrow$  localTime
2: if  $\mathcal{G} \in \{\mathcal{G}_{\text{IMPERFLCLOCK}}, \mathcal{G}_{\text{RO}}\}$  then
3:   if  $\mathcal{G} = \mathcal{G}_{\text{IMPERFLCLOCK}}$  then set isSync  $\leftarrow$  false
4:   Send (DE-REGISTER, sid) to  $\mathcal{G}$  and set registration status as de-registered with  $\mathcal{G}$ .
5:   Output the valued received by  $\mathcal{G}$ .
   end if
6: if  $\mathcal{G} = \mathcal{G}_{\text{LEDGER}}$  then
7:   Set isSync  $\leftarrow$  false
   Send (DE-REGISTER, sid) to  $\mathcal{F}_{\text{N-MC}}^{\Delta}$ , set its registration status as de-registered with  $\mathcal{F}_{\text{N-MC}}^{\Delta}$  and output
   (DE-REGISTER, sid, P).
   end if

```

### C.3 Ledger Maintenance

What exactly a party might execute in a round depends on its status: newly registered parties first run through initialization and only later start to create blocks. All these steps are grouped in the main ledger operation procedure below and the relevant processes follow in the following paragraph. Compared to Genesis, the main round procedure is substantially more involved and we try to modularly adjust the code to see the difference.

**Protocol LedgerMaintenance( $\mathcal{C}_{\text{loc}}, P, \text{sid}, k, s, R, f$ )**

The following steps are executed in an (MAINTAIN-LEDGER, sid, minerID)-interruptible manner:

```

1: if isInit is false then invoke Initialization-Chronos(P, sid, R); if Initialization-Chronos halts then halt (this will
   abort the execution)
   end if
2: // From here the variables  $v_p^{\text{vrf}}, v_p^{\text{kes}}, \text{localTime}, \text{ep}, \text{s1}, \mathcal{C}_{\text{loc}}, \text{isSync}, T_p^{\text{ep}}, T_p^{\text{ep, bc}}, \text{fetchCompleted}, t_{\text{work}}$  can be
   used to read from as they are guaranteed to be initialized.
3: if isSync and stalled before (and now up and running) then
4:   SimulateClockAdjustments(P, R, k, f, s)
5: end if
6: if not isSync then
7:   Call JoinProc(P, sid, R, k, f, s,  $t_{\text{off}}, t_{\text{stable}}, t_{\text{minSync}}$ )
8: end if
9: // normal operation when alert
10: Call FetchInformation(P, sid) and denote the output by  $(\mathcal{C}_1, \dots, \mathcal{C}_M), (\text{tx}_1, \dots, \text{tx}_k)$ .
11: Set buffer  $\leftarrow$  buffer  $\parallel$   $(\text{tx}_1, \dots, \text{tx}_k)$  and define futureChains  $\leftarrow$  futureChains  $\parallel$   $(\mathcal{C}_1, \dots, \mathcal{C}_M)$ 
12: Call UpdateTime(P, R)
13: // Ensures the processing of new beacons arrived in chains only.
14: Extract beacons  $\mathcal{B} \leftarrow \{\text{SB}_1, \dots, \text{SB}_n\}$  contained in  $\mathcal{C}_1, \dots, \mathcal{C}_M$  and not yet contained in syncBuffer.
15: Call ProcessBeacons(P, sid,  $\mathcal{B}$ )
16: Let  $\mathcal{N}_0$  be the subsequence of futureChains s.t.  $\mathcal{C} \in \mathcal{N}_0 : \Leftrightarrow \forall B \in \mathcal{C} : \text{slotnum}(B) \leq \text{localTime}$ 

```



```

17: Remove each  $\mathcal{C} \in \mathcal{N}_0$  from futureChains.
18: fetchCompleted  $\leftarrow$  true
19: Call SelectChain( $P, \text{sid}, \mathcal{C}_{\text{loc}}, \mathcal{N}_0, k, s, R, f$ ) to update  $\mathcal{C}_{\text{loc}}$ 
20: if  $t_{\text{work}} < \text{localTime}$  then
21:   Call UpdateStakeDist( $P, \text{sid}, k, R, f$ ) to update the values  $\mathbb{S}_{\text{ep}}, \alpha_P^{\text{ep}}, T_P^{\text{ep}}, T_P^{\text{ep, bc}}$  and  $\eta_{\text{ep}}$ .
22:   Call StakingProcedure( $k, P, \text{ep}, \text{s1}, \text{buffer}, \mathcal{C}_{\text{loc}}$ )
23:   Set  $t_{\text{work}} \leftarrow \text{localTime}$ 
24:   if  $\text{localTime} \bmod R = 0$  then
25:     Call SyncProc( $P, \text{sid}, R$ )
26:   end if
  end if
27: Call FinishRound( $P$ ) // Mark normal round actions as finished.

```

**Standard Round Operations.** Before discussing the involved special case of joining, let us first define all relevant procedures for an alert party to complete its task per round as described in Section 5 of the main body, that is:

- Fetch information from the network (by a call to `FetchInformation`).
- Update the time (by a call to `UpdateTime`): the party locally advances its time-stamp whenever it realizes that a new round has started by a call  $\mathcal{G}_{\text{IMPERFLOCK}}$  and comparing it to `lastTick`.
- Record the arrival times of the synchronization beacons the protocol sends out (call to `ProcessBeacons`).
- Process the received chains: as some chains might be created by parties whose time-stamps might be ahead, the future chains are stored in the buffer `futureChains` for later usage. Among the remaining chains, the protocol will according to the Genesis chain-selection rule decide whether any chain is more preferable than the local chain (procedure `SelectChain`).
- Run the main staking procedure (`StakingProcedure`) to evaluate slot leadership, and potentially create and emit a new block or synchronization beacon. Before the main staking procedure is executed, the local state is updated including the current stake distribution (call to `UpdateStakeDist`).
- If the end of the round coincides with the end of an epoch, the synchronization procedure is executed. This core procedure of our proposal is detailed below.

The above round operations are based on several core procedures of blockchain protocols, most notably three procedures we introduce next:

- The initialization algorithm to properly define the initial values.
- The verification algorithms to evaluate whether a chain is valid, and for the specific case of Ouroboros Chronos also whether a beacon is valid.

#### C.4 Initialization

The first special procedure a party runs through is initialization. It is invoked upon the first MAINTAIN-LEDGER input given to this instance. Since every party starts at local time 0 and has no knowledge whether the session is already running, it will first try to claim stake from  $\mathcal{F}_{\text{INIT}}$  in its first round and then in subsequent rounds, tries to receive the genesis block until it is successful. As a difference to Ouroboros Genesis, a party might be delayed in receiving the genesis block. In any case, once a party obtains the genesis block it will initialize the variables of this instance which are described in Table M.1. In particular, as specified by our setup  $\mathcal{F}_{\text{INIT}}$  if the genesis block is delivered to an inaugural party of the system, then the party considers itself as synchronized.<sup>19</sup> Otherwise it has to invoke the joining procedure.

<sup>19</sup> Note that this knowledge is needed to bootstrap the system with a set of alert parties.

**Protocol Initialization-Chronos( $P, sid, R$ )**

The following steps are executed in an (MAINTAIN-LEDGER,  $sid$ ,  $minerID$ )-interruptible manner:

- 1: Send (KeyGen,  $sid, P$ ) to  $\mathcal{F}_{VRF}$  and  $\mathcal{F}_{KES}$ ; receiving (VerificationKey,  $sid, v_p^{vrf}$ ) and (VerificationKey,  $sid, v_p^{kes}$ ), respectively. // Note that this can be seen as defining the party's address  $\langle v_p^{vrf}, v_p^{kes} \rangle$  (where  $\langle \rangle$  denotes an encoding).
- 2: **if**  $localTime = 0$  **then** // Right after registration of this instance
- 3:     Send (ver\_keys,  $sid, P, v_p^{vrf}, v_p^{kes}$ ) to  $\mathcal{F}_{INIT}$  to claim stake from the genesis block.
- 4:     FinishRound( $P$ ) // Mark round actions as finished. Resume below upon next activation
- 5:     Call UpdateTime( $P, R, f$ ) to update  $localTime, ep$ , and  $s1$
- 6:     **while**  $localTime = 0$  **do**
- 7:         Call UpdateTime( $P, R, f$ ) to update  $localTime, ep$ , and  $s1$  and give up the activation (set anchor here)
- 8:     **end while**
- 9: **end if**  
    // The following is executed in future init-activations of this instance
- 10: **if**  $localTime > 0$  **then**
- 11:     **if**  $\mathcal{F}_{INIT}$  signals an error **then**
- 12:         Halt the execution.
- 13:     **end if**
- 14:     Send (genblock\_req,  $sid, P$ ) to  $\mathcal{F}_{INIT}$ .
- 15:     **while**  $\mathcal{F}_{INIT}$  ignores the input **do**
- 16:         FinishRound( $P$ ) // Round actions as finished. Resume below upon next activation
- 17:         Send (genblock\_req,  $sid, P$ ) to  $\mathcal{F}_{INIT}$ .
- 18:     **end while**
- 19:     Receive the genesis block (genblock,  $sid, \mathbf{G} = (\mathcal{S}_1, \eta_1)$ ), where
 
$$\mathcal{S}_1 = \left( (U_1, v_1^{vrf}, v_1^{kes}, s_1), \dots, (U_n, v_n^{vrf}, v_n^{kes}, s_n) \right) .$$
- 20:     Set  $\mathcal{C}_{loc} \leftarrow (\mathbf{G})$ .
- 21:     **if**  $\mathcal{F}_{INIT}$  did not mark the returned value as Running (the execution just started) and this instance was time-aware and online from registration onward, then additionally set  $isSync \leftarrow true$  // This party becomes a synchronized inaugural protocol participant.
- 22:     **end if**
- 23:     Set  $isInit \leftarrow true$ ,  $fetchCompleted \leftarrow false$ ,  $t_{work} \leftarrow 0$ ,  $lastTimeAlert \leftarrow 0$ ,  $localTime, s1, ep \leftarrow 1$  and call UpdateStakeDist( $P, sid, k, R, f$ ) to initialize the values  $S_{ep}, \alpha_p^{ep}, T_p^{ep}, T_p^{ep, bc}$  and  $\eta_{ep}$ . // Now, ready to start, either with the first slot in first epoch, or with bootstrapping (joining procedure).
- 24:     **buffer**  $\leftarrow \emptyset$ , **futureChains**  $\leftarrow \emptyset$
- 25:     EpochUpdate( $\cdot$ )  $\leftarrow$  empty table (initial symbol  $\perp$ ), EpochUpdate(0)  $\leftarrow$  Done

### C.5 Validity Checks of Chains and Beacons in Chronos

**Chain verification.** A core procedure is to distinguish valid from invalid blockchains. The procedure is depicted below. It is a slightly extended version of Ouroboros Genesis to cover beacons.

**Protocol IsValidChain( $P, sid, \mathcal{C}, f, R$ )**

- if**  $\mathcal{C}$  contains empty epochs or starts with a block other than  $\mathbf{G}$ , or  $isvalidstate(st) = 0$  **then**  
     **return false**  
**end if**  
**if**  $isSync$  and  $(\exists B \in \mathcal{C} : slotnum(B) > localTime)$  **then**  
     **return false**  
**end if**  
**for** each epoch  $ep$  **do**

```

// Derive stake distribution and randomness for this epoch from  $\mathcal{C}$ 
// In the following,  $H(\cdot)$  stands for an RO evaluation for simplicity.
Set  $\mathbb{S}_{\text{ep}}^{\mathcal{C}}$  to be the stakeholder distribution at the end of epoch  $\text{ep} - 2$  in  $\mathcal{C}$ .
Set  $\alpha_{\mathcal{P}'}^{\text{ep}, \mathcal{C}}$  to be the relative stake of any party  $\mathcal{P}'$  in  $\mathbb{S}_{\text{ep}}^{\mathcal{C}}$  and  $T_{\mathcal{P}'}^{\text{ep}, \mathcal{C}} \leftarrow 2^{\ell_{\text{VRF}}} \phi_f(\alpha_{\mathcal{P}'}^{\text{ep}, \mathcal{C}})$ . Also define
 $T_{\mathcal{P}'}^{\text{ep}, \text{bc}, \mathcal{C}} \leftarrow 2^{\ell_{\text{VRF}}} \cdot \alpha_{\mathcal{P}'}^{\text{ep}, \mathcal{C}}$ .
Set  $\eta_{\text{ep}}^{\mathcal{C}} \leftarrow H(\eta_{\text{ep}-1}^{\mathcal{C}} \parallel \text{ep} \parallel v)$  where  $v$  is the concatenation of the VRF outputs  $y_{\rho}$  from all blocks in  $\mathcal{C}$  from
the first two-thirds of slots of epoch  $\text{ep} - 1$ , and  $\eta_1^{\mathcal{C}} \triangleq \eta_1$  from  $\mathbf{G}$ .
for each block  $B$  in  $\mathcal{C}$  from epoch  $\text{ep}$  do
  Parse  $B$  as  $(h, \text{st}, \text{s1}, \text{crt}, \rho, \sigma)$ .
  // Check hash
  Set  $\text{badhash} \leftarrow (h \neq H(B^{-1}))$ , where  $B^{-1}$  is the last block in  $\mathcal{C}$  before  $B$ .
  // Check VRF values
  Parse  $\text{crt}$  as  $(\mathcal{P}', y, \pi)$  for some  $\mathcal{P}'$ .
  Send  $(\text{Verify}, \text{sid}, \eta_{\text{ep}} \parallel \text{s1} \parallel \text{TEST}, y, \pi, v_{\mathcal{P}'}^{\text{vrf}})$  to  $\mathcal{F}_{\text{VRF}}$ ,
    denote its response by  $(\text{Verified}, \text{sid}, \eta_{\text{ep}} \parallel \text{s1} \parallel \text{TEST}, y, \pi, b_1)$ .
  Send  $(\text{Verify}, \text{sid}, \eta_{\text{ep}} \parallel \text{s1} \parallel \text{NONCE}, y_{\rho}, \pi_{\rho}, v_{\mathcal{P}'}^{\text{vrf}})$  to  $\mathcal{F}_{\text{VRF}}$ ,
    denote its response by  $(\text{Verified}, \text{sid}, \eta_{\text{ep}} \parallel \text{s1} \parallel \text{NONCE}, y_{\rho}, \pi_{\rho}, b_2)$ ,
  Set  $\text{badvrf} \leftarrow (b_1 = 0 \vee b_2 = 0 \vee y \geq T_{\mathcal{P}'}^{\text{ep}, \mathcal{C}})$ .
  // Check signature
  Send  $(\text{Verify}, \text{sid}, (h, \text{st}, \text{s1}, \text{crt}, \rho), \text{s1}, \sigma, v_{\mathcal{P}'}^{\text{kes}})$  to  $\mathcal{F}_{\text{KES}}$ ,
    denote its response by  $(\text{Verified}, \text{sid}, (h, \text{st}, \text{s1}, \text{crt}, \rho), \text{s1}, b_3)$ .
  Set  $\text{badsig} \leftarrow (b_3 = 0)$ .
  // Check Beacons
  if  $\exists \text{SB} \in B \wedge \text{slotnum}(B) > (\text{ep} - 1)R + 2R/3$  then
    Set  $\text{badBeacon} \leftarrow \text{true}$ 
  else if  $\exists \text{SB} \in B : \text{slotnum}(\text{SB}) > \text{slotnum}(B) \vee \text{slotnum}(\text{SB}) \notin [(\text{ep} - 1)R + 1, \text{ep} \cdot R]$  then
    Set  $\text{badBeacon} \leftarrow \text{true}$ 
  else
    for each  $\text{SB} \in B$  do
      Parse  $\text{SB}$  as  $(\text{s1}', \mathcal{P}', y, \pi)$ 
      If  $\mathcal{C}$  contains more than one beacon with  $(\text{s1}', \mathcal{P}', \cdot, \cdot)$  then set  $\text{badBeacon} \leftarrow \text{true}$ 
      Send  $(\text{Verify}, \text{sid}, \eta_{\text{ep}'} \parallel \text{s1}' \parallel \text{SYNC}, y, \pi, v_{\mathcal{P}'}^{\text{vrf}})$  to  $\mathcal{F}_{\text{VRF}}$ .
        Denote the response from  $\mathcal{F}_{\text{VRF}}$  by  $(\text{Verified}, \text{sid}, \eta_{\text{ep}'} \parallel \text{s1}' \parallel \text{SYNC}, y, \pi, b_4)$ ,
      if  $(b_4 = 0)$  or  $(y \geq T_{\mathcal{P}'}^{\text{ep}, \text{bc}, \mathcal{C}})$  then
        Set  $\text{badBeacon} \leftarrow \text{true}$ 
      end if
    end for
  end if
  if  $(\text{badhash} \vee \text{badvrf} \vee \text{badsig} \vee \text{badBeacon})$  then
    return false
  end if
end for
return truek

```

**The beacon validity predicate.** Beacons validity is related to chain validity as one has to verify validity of leadership. The details are found below:

**Protocol** ValidSB( $\mathcal{P}, \text{sid}, \text{SB}, \mathcal{C}, f, R$ )

// Precondition: Chain  $\mathcal{C}$  is valid. Returns true if the beacon is a valid beacon w.r.t.  $\mathcal{C}$ , undecided if no judgement is possible, and false if the beacon is invalid w.r.t.  $\mathcal{C}$ .

```

Parse SB as  $(\mathbf{s}1', P', y, \pi)$ 
Let  $ep'$  be the epoch number slot  $\mathbf{s}1'$  falls into. Let  $ep := ep' - 2$ .
if  $\mathcal{C}$  contains no block in epoch  $ep'$  then
    return undecided // no judgement possible for this beacon
end if
// Derive stake distribution and randomness for epoch  $ep'$  as indicated by  $\mathcal{C}$ 
Set  $\mathbb{S}_{ep'}^{\mathcal{C}}$  to be the stakeholder distribution at the end of epoch  $ep' - 2$  in  $\mathcal{C}$ .
Set  $\alpha_{P'}^{ep', \mathcal{C}}$  to be the relative stake of party  $P'$  in  $\mathbb{S}_{ep'}^{\mathcal{C}}$  and  $T_{P'}^{ep', bc, \mathcal{C}} \leftarrow 2^{\ell_{\text{VRF}}} \cdot \alpha_{P'}^{ep', \mathcal{C}}$ .
Set  $\eta_{ep'}^{\mathcal{C}} \leftarrow H(\eta_{ep'-1}^{\mathcal{C}} \parallel ep' \parallel v)$  where  $v$  is the concatenation of the VRF outputs  $y_\rho$  from the existing blocks in  $\mathcal{C}$  with slot numbers of the first two-thirds slots of epoch  $ep' - 1$  (and  $\eta_1^{\mathcal{C}} \triangleq \eta_1$  from  $\mathbf{G}$ ).
// Check VRF value
Send (Verify,  $sid, \eta_{ep'} \parallel \mathbf{s}1' \parallel \text{SYNC}, y, \pi, v_{P'}^{\text{VRF}}$ ) to  $\mathcal{F}_{\text{VRF}}$ .
Denote the response from  $\mathcal{F}_{\text{VRF}}$  by (Verified,  $sid, \eta_{ep'} \parallel \mathbf{s}1' \parallel \text{SYNC}, y, \pi, b_1$ ),
if  $b_1 = 0$  or  $y \geq T_{P'}^{ep', bc, \mathcal{C}}$  then
    return false
end if
return true

```

## C.6 Fetching information, stake distribution and time update

The two algorithms FetchInformation and UpdateTime are basically identical to Ouroboros Genesis except that FetchInformation was simplified since newly joining parties do not make an active request.

### Protocol FetchInformation( $P, sid$ )

```

1: if fetchCompleted then
2:   Set fetchcount  $\leftarrow 0$ 
3: else
4:   Set fetchcount := 1 // Compared to Genesis, time-aware and online parties in Chronos do always fetch once
   per round and never have to catch up missed round messages.
5: end if
// Fetching on  $\mathcal{F}_{N\text{-MC}}^{bc}$ .
6: Send fetchcount fetch-queries (FETCH, sid) to  $\mathcal{F}_{N\text{-MC}}^{bc}$ ; denote the  $i$ th response from  $\mathcal{F}_{N\text{-MC}}^{bc}$  by (FETCH, sid,  $b_i$ ).
7: Extract chains  $\mathcal{C}_1, \dots, \mathcal{C}_k$  from  $b_1 \dots b_{\text{fetchcount}}$ .
// Fetching on  $\mathcal{F}_{N\text{-MC}}^{tx}$ .
8: Send fetchcount fetch-queries (FETCH, sid) to  $\mathcal{F}_{N\text{-MC}}^{tx}$ ; denote the  $i$ th response from  $\mathcal{F}_{N\text{-MC}}^{tx}$  by (FETCH, sid,  $b_i$ ).
9: Extract received transactions  $(\mathbf{tx}_1, \dots, \mathbf{tx}_k)$  from  $b_1 \dots b_{\text{fetchcount}}$ .
10: if not isSync or  $P$  is stalled then
11:   buffer  $\leftarrow$  buffer ||  $(\mathbf{tx}_1, \dots, \mathbf{tx}_n)$ 
12:   futureChains  $\leftarrow$  futureChains  $\cup \{\mathcal{C}_1, \dots, \mathcal{C}_n\}$ 
13: end if

```

OUTPUT: The protocol outputs  $(\mathcal{C}_1, \dots, \mathcal{C}_k)$  and  $(\mathbf{tx}_1, \dots, \mathbf{tx}_k)$  to its caller (but not to  $\mathcal{Z}$ ).

The stake distributions for epochs defined in the local chain (and all associated state-variables) are computed as follows:

### Protocol UpdateStakeDist( $P, k, R, f$ )

```

1: Set  $\mathbb{S}_{ep}$  to be the stakeholder distribution at the end of epoch  $ep - 2$  in  $\mathcal{C}_{\text{loc}}$  in case  $ep \geq 2$  (and keep the initial
   stake distribution in case  $ep < 2$ ).
2: Set  $\alpha_P^{ep}$  to be the relative stake of  $P$  in  $\mathbb{S}_{ep}$  and  $T_P^{ep} \leftarrow 2^{\ell_{\text{VRF}}} \phi_f(\alpha_P^{ep})$  as well as  $T_P^{ep, bc} \leftarrow 2^{\ell_{\text{VRF}}} \cdot \alpha_P^{ep}$ 

```

3: Set  $\eta_{ep} \leftarrow H(\eta_{ep-1} \parallel \mathbf{ep} \parallel v)$  where  $v$  is the concatenation of the VRF outputs  $y_\rho$  from all blocks in  $C_{loc}$  from the first  $2R/3$  slots of epoch  $\mathbf{ep} - 1$  (and if  $\mathbf{ep} = 1$ ,  $\eta_1$  is the nonce of the genesis block).  
 OUTPUT: The protocol outputs  $\mathbb{S}_{ep}, \alpha_P^{ep}, T_P^{ep}, \eta_{ep}$  and  $T_P^{ep, bc}$  to its caller (but not to  $\mathcal{Z}$ ).

And, finally, the time update procedure:

#### Protocol UpdateTime( $P, R$ )

// Precondition: Only executed if time-aware  
 1: Send (CLOCK-GET,  $\text{sid}_C$ ) to  $\mathcal{G}_{\text{IMPERFLOCK}}$  and receive (CLOCK-GET,  $\text{sid}_C$ , tick)  
 2: **if** lastTick  $\neq$  tick **then**  
 3:   lastTick  $\leftarrow$  tick  
 4:   localTime  $\leftarrow$  localTime + 1  
 5:   fetchCompleted  $\leftarrow$  false  
 6: **end if**  
 7: Set  $\mathbf{ep} \leftarrow \lceil \text{localTime}/R \rceil$ , and  $\mathbf{s1} \leftarrow \text{localTime}$ .  
 OUTPUT: The protocol outputs localTime, ep, s1 to its caller (but not to  $\mathcal{Z}$ ).

## C.7 Process Beacons and Arrival Times

The following procedure records and processes beacons, their arrival times, and filters out invalid beacons. Special care needs to be made to properly filter out bogus beacons as soon as possible. The validity predicate for beacons follows in the next section, where we discuss all validity predicates in Ouroboros Chronos.

#### Protocol ProcessBeacons( $P, \text{sid}, \mathcal{B}$ )

1: **if not** fetchCompleted **then**  
 2:   Send (FETCH, sid) to  $\mathcal{F}_{N-MC}^{\text{sync}}$ . denote the  $i$ th response from  $\mathcal{F}_{N-MC}^{\text{sync}}$  by (FETCH, sid,  $b$ ).  
 3:   Extract all received beacons ( $\text{SB}_1, \dots, \text{SB}_k$ ) contained in  $b \cup \mathcal{B}$ .  
 4:   **for each**  $\text{SB}_i$  with  $\text{Timestamp}_{\text{SB}}(\text{SB}_i) = \perp$  **do**  
 5:     syncBuffer  $\leftarrow$  syncBuffer  $\cup$  { $\text{SB}_i$ }  
 6:     Let  $ep$  be the epoch number slotnum( $\text{SB}_i$ ) belongs to  
 7:     **if** isSync  $\wedge$  (EpochUpdate( $ep - 1$ ) = Done) **then**  
 8:       Set  $\text{Timestamp}_{\text{SB}}(\text{SB}_i) \leftarrow (\text{localTime}, \text{final})$ . // The measurement is final.  
 9:     **else**  
 10:        $\text{Timestamp}_{\text{SB}}(\text{SB}_i) \leftarrow (\text{localTime}, \text{temp})$  // Will be adjusted upon next time shift.  
 11:     **end if**  
 12:   **end for**  
 13:   // Buffer cleaning. Keep one representative arrival time.  
 14:   **if** isSync **then**  
 15:     Remove from syncBuffer all beacons such that ValidSB( $P, \text{sid}, \text{SB}, C_{loc}, f, R$ ) returns false.  
 16:     syncBuffer<sub>valid</sub>  $\leftarrow$  { $\text{SB}' \in \text{syncBuffer} \mid \text{ValidSB}(P, \text{sid}, \text{SB}', C_{loc}, f, R) = \text{true}$ }  
 17:     Let  $L = (\text{SB}_1, \dots, \text{SB}_n)$  be a canonical ordering of syncBuffer<sub>valid</sub>  
 18:     **for each**  $\text{SB} = (\mathbf{s1}, P, y, \pi) \in L$  **do**  
 19:        $Q_{\text{SB}} \leftarrow \{\text{SB}' = (\mathbf{s1}', P', \cdot, \cdot) \in L \mid P' = P \wedge \mathbf{s1}' = \mathbf{s1}\}$   
 20:        $\text{min}_{\text{SB}} \leftarrow \min\{\text{Timestamp}(\text{SB}') \mid \text{SB}' \in Q_{\text{SB}}\}$   
 21:        $\text{SB}' \leftarrow \min\{\text{SB}'' \in Q_{\text{SB}} \mid \text{Timestamp}(\text{SB}'') = \text{min}_{\text{SB}}\}$  // Min w.r.t. ordering in  $L$   
 22:       Remove from syncBuffer all beacons  $(\mathbf{s1}, P, \cdot, \cdot)$  except  $\text{SB}'$ .  
 23:     **end for**  
 24:   **end if**  
 25: **end if**  
 OUTPUT: ok to its caller (but not to  $\mathcal{Z}$ ).

## C.8 Select Chain

Chain selection consists of two steps: filtering out valid chains, and second compare them using the Genesis rule.

**Protocol**  $\text{SelectChain}(\mathcal{P}, \text{sid}, \mathcal{C}_{\text{loc}}, \mathcal{N} = \{\mathcal{C}_1, \dots, \mathcal{C}_M\}, k, s, R, f)$

```

1: Initialize  $\mathcal{N}_{\text{valid}} \leftarrow \emptyset$ 
2: for  $i = 1 \dots M$  do
3:   Invoke  $\text{IsValidChain}(\mathcal{P}, \text{sid}, \mathcal{C}_i, f, R)$ ; if it returns true then update  $\mathcal{N}_{\text{valid}} \leftarrow \mathcal{N}_{\text{valid}} \cup \mathcal{C}_i$ 
   end for
4: Execute Algorithm  $\text{maxvalid-bg}(\mathcal{C}_{\text{loc}}, \mathcal{N}_{\text{valid}} = \{\mathcal{C}_1, \dots, \mathcal{C}_M\}, k, s, f)$  and receive its output  $\mathcal{C}_{\text{max}}$ .
5: Replace  $\mathcal{C}_{\text{loc}}$  by  $\mathcal{C}_{\text{max}}$ 
OUTPUT: The protocol outputs  $\mathcal{C}_{\text{max}}$  to its caller (but not to  $\mathcal{Z}$ ).

```

## C.9 The Genesis Chain Selection Rule

The genesis rule is given below:

**Algorithm**  $\text{maxvalid-bg}(\mathcal{C}_{\text{loc}}, \mathcal{N} = \{\mathcal{C}_1, \dots, \mathcal{C}_M\}, k, s, f)$

```

// Compare  $\mathcal{C}_{\text{max}}$  to each  $\mathcal{C}_i \in \mathcal{N}$ 
1: Set  $\mathcal{C}_{\text{max}} \leftarrow \mathcal{C}_{\text{loc}}$ .
2: for  $i = 1$  to  $M$  do
3:   if ( $\mathcal{C}_i$  forks from  $\mathcal{C}_{\text{max}}$  at most  $k$  blocks) then
4:     if  $|\mathcal{C}_i| > |\mathcal{C}_{\text{max}}|$  then // Condition A
       Set  $\mathcal{C}_{\text{max}} \leftarrow \mathcal{C}_i$ .
     end if
5:   else
6:     Let  $j \leftarrow \max \{j' \geq 0 \mid \mathcal{C}_{\text{max}} \text{ and } \mathcal{C}_i \text{ have the same block in } \text{s1}_{j'}\}$ 
7:     if  $|\mathcal{C}_i[0 : j + s]| > |\mathcal{C}_{\text{max}}[0 : j + s]|$  then // Condition B
       Set  $\mathcal{C}_{\text{max}} \leftarrow \mathcal{C}_i$ .
     end if
   end if
end for
8: return  $\mathcal{C}_{\text{max}}$ .

```

## C.10 The Core Staking Procedure of Alert Parties

Once a party has properly filtered all information and updated its state, it can run the core staking procedure and beacon emitting process formally given below:

**Protocol**  $\text{StakingProcedure}(\mathcal{P}, \text{sid}, k, \text{ep}, \text{s1}, \text{buffer}, \mathcal{C}_{\text{loc}})$

The following steps are executed in an (MAINTAIN-LEDGER, sid, minerID)-interruptible manner:

```

// Determine leader status
1: Send (EvalProve, sid,  $\eta_j \parallel \text{s1} \parallel \text{NONCE}$ ) to  $\mathcal{F}_{\text{VRF}}$ , denote the response from  $\mathcal{F}_{\text{VRF}}$  by (Evaluated, sid,  $y_\rho, \pi_\rho$ ).
2: Send (EvalProve, sid,  $\eta_j \parallel \text{s1} \parallel \text{SYNC}$ ) to  $\mathcal{F}_{\text{VRF}}$ , denote the response from  $\mathcal{F}_{\text{VRF}}$  by (Evaluated, sid,  $y_{\text{sync}}, \pi_{\text{sync}}$ ).
3: Send (EvalProve, sid,  $\eta_j \parallel \text{s1} \parallel \text{TEST}$ ) to  $\mathcal{F}_{\text{VRF}}$ , denote the response from  $\mathcal{F}_{\text{VRF}}$  by (Evaluated, sid,  $y, \pi$ ).
4: if  $y < T_{\text{p}}^{\text{ep}}$  and this party is sign-capable then
   // Generate a new block

```

```

5: Set  $\text{buffer}' \leftarrow \text{buffer}$ ,  $N \leftarrow \text{tx}_P^{\text{base-tx}}$ , and  $\text{st} \leftarrow \text{blockify}_{\text{OC}}(N)$ 
6: repeat
7:   Parse  $\text{buffer}'$  as sequence  $(\text{tx}_1, \dots, \text{tx}_n)$ 
8:   for  $i = 1$  to  $n$  do
9:     if  $\text{ValidT}_{\text{XOC}}(\text{tx}_i, \text{st} || \text{st}) = 1$  then
10:       $N \leftarrow N || \text{tx}_i$ 
11:      Remove  $\text{tx}$  from  $\text{buffer}'$ 
12:      Set  $\text{st} \leftarrow \text{blockify}_{\text{OC}}(N)$ 
    end if
  end for
  until  $N$  does not increase anymore
13: Set  $\text{crt} = (P, y, \pi)$ ,  $\rho = (y_\rho, \pi_\rho)$  and  $h \leftarrow H(\text{head}(\mathcal{C}_{\text{loc}}))$ .
14: if Slot  $\text{s1}$  is within the first  $2R/3$  slots of this epoch then
15:    $\text{sb} \leftarrow \{\text{SB}' \in \text{syncBuffer} \mid \text{ValidSB}(P, \text{sid}, \text{SB}', \mathcal{C}_{\text{loc}}, f, R) = \text{true}\}$ 
16:   Remove from  $\text{sb}$  all beacons  $\text{SB} = (\text{s1}, P, \cdot, \cdot)$  that satisfy:
17:    $(\text{slotnum}(\text{SB}) > \text{s1}) \vee (\text{slotnum}(\text{SB}) \leq (\text{ep} - 1) \cdot R) \vee \mathcal{C}_{\text{loc}}$  contains a beacon  $(\text{s1}, P, \cdot, \cdot)$ 
18: end if
19: Send  $(\text{USign}, \text{sid}, P, (h, \text{st}, \text{sb}, \text{s1}, \text{crt}, \rho), \text{s1})$  to  $\mathcal{F}_{\text{KES}}$ ; denote the response from  $\mathcal{F}_{\text{KES}}$  by  $(\text{Signature}, \text{sid}, (h, \text{st}, \text{sb}, \text{s1}, \text{crt}, \rho), \text{s1}, \sigma)$ .
20: Set  $B \leftarrow (h, \text{st}, \text{sb}, \text{s1}, \text{crt}, \rho, \sigma)$  and update  $\mathcal{C}_{\text{loc}} \leftarrow \mathcal{C}_{\text{loc}} || B$ .
  // Multicast the extended chain and wait.
21: Send  $(\text{MULTICAST}, \text{sid}, \mathcal{C}_{\text{loc}})$  to  $\mathcal{F}_{\text{N-MC}}^{\text{bc}}$  and proceed from here upon next activation of this procedure.
22: else
23:   Evolve the KES signing key at least to  $\text{localTime}$  by sending  $(\text{USign}, \text{sid}, P, 0, \text{s1})$  to  $\mathcal{F}_{\text{KES}}$  (and ignore the returned value). Give up activation and set anchor here to resume on next maintenance activation
  end if
24: if  $y_{\text{sync}} < T_P^{\text{ep, bc}}$  and  $\text{s1}$  lies within the first  $R/6$  slots of this epoch then
25:    $\text{SB} \leftarrow (\text{s1}, P, y_{\text{sync}}, \pi_{\text{sync}})$ .
26:   Send  $(\text{MULTICAST}, \text{sid}, \text{SB})$  to  $\mathcal{F}_{\text{N-MC}}^{\text{sync}}$  and set anchor at end of procedure to resume on next maintenance activation
27: else
28:   Give up activation and set anchor at end of procedure to resume on next maintenance activation
29: end if

```

### C.11 Code of the Synchronization Procedure

The synchronization procedure is run at an epoch boundary, and the code is given below:

#### Protocol SyncProc( $P, \text{sid}, R$ )

```

1: // Only called when:  $P$  is alert,  $\text{localTime} \bmod R = 0$  and  $\text{localTime} > 0$ 
2: Set  $i \leftarrow \text{localTime} \text{ div } R$ 
3: if (not  $\text{EpochUpdate}(i) = \text{Done}$ ) then
4:    $\text{EpochUpdate}(i) \leftarrow \text{Done}$  // Remember that clock adjustment has happened
5:    $\mathcal{B}_i \leftarrow \mathcal{C}_{\text{loc}}[(i-1)R : (i-1)R + 2R/3]$ 
6:    $\mathcal{S}_i \leftarrow \{\text{SB} \mid \exists B \in \mathcal{B}_i : \text{SB} \in B \wedge \text{slotnum}(\text{SB}) \in \{(i-1)R, \dots, (i-1)R + R/6\}\}$ 
7:   for each  $\text{SB} = (\text{s1}, P, y, \pi) \in \mathcal{S}_i$  do
8:     // Find representative beacon and compute recommendation.
9:     Find unique  $\text{SB}' = (\text{s1}, P, \cdot, \cdot) \in \text{syncBuffer}$ . If inexistent, set  $\text{SB}' \leftarrow \perp$ .
10:    if  $\text{SB}' \neq \perp$  then
11:      Set  $\text{Timestamp}_{\text{SB}}(\text{SB}) \leftarrow \text{Timestamp}_{\text{SB}}(\text{SB}')$  // Assign correct value
12:       $\text{recom}(\text{SB}) \leftarrow \text{slotnum}(\text{SB}) - \text{Timestamp}(\text{SB})$ 
13:    else

```

```

14:      $S \leftarrow S \setminus \{\text{SB}\}$  // Negligible probability event in execution.
15:   end if
16: end for
17:  $\text{shift}_i \leftarrow \text{med} \{\text{recom}(\text{SB}) \mid \text{SB} \in \mathcal{S}_i\}$ 
18: for each SB with  $\text{Timestamp}_{\text{SB}}(\text{SB}) = (a, \text{temp})$  do
19:    $\text{Timestamp}_{\text{SB}}(\text{SB}) \leftarrow (a + \text{shift}_i, \text{final})$ 
20: end for
21: if  $\text{shift}_i > 0$  then // Move fast forward
22:    $\text{newTime} \leftarrow \text{localTime} + \text{shift}_i$ 
23:    $\mathcal{M}_{\text{chains}} \leftarrow \mathcal{M}_{\text{sync}} \leftarrow \emptyset$ 
24:   while  $\text{localTime} < \text{newTime}$  do
25:      $\text{localTime} \leftarrow \text{localTime} + 1$ 
26:     Let  $\mathcal{N}_0$  be the subsequence of  $\text{futureChains}$  s.t.
27:        $\mathcal{C} \in \mathcal{N}_0 \Leftrightarrow \forall B \in \mathcal{C} : \text{slotnum}(B) \leq \text{localTime}$ 
28:     Remove each  $\mathcal{C} \in \mathcal{N}_0$  from  $\text{futureChains}$ .
29:     Call  $\text{SelectChain}(\mathcal{C}_{\text{loc}}, \mathcal{N}_0, k, s, R, f)$  to update  $\mathcal{C}_{\text{loc}}$ 
30:     Call  $\text{UpdateStakeDist}(P, k, R, f)$ 
31:     Emulate  $\text{StakingProcedure}(k, P, \text{ep}, \text{sl}, \text{buffer}, \mathcal{C}_{\text{loc}})$  but instead of multicasting new chains or
    beacons, add them to the sets  $\mathcal{M}_{\text{chains}}$  and  $\mathcal{M}_{\text{sync}}$ , respectively (activation is not lost).
32:   end while
33:   Send (MULTICAST, sid,  $\mathcal{M}_{\text{chains}}$ ) to  $\mathcal{F}_{\text{N-MC}}^{\text{bc}}$  and (MULTICAST, sid,  $\mathcal{M}_{\text{sync}}$ ) to  $\mathcal{F}_{\text{N-MC}}^{\text{sync}}$  and proceed from here
    upon next activation of this procedure.
34: end if
35: if  $\text{shift}_i < 0$  then // Need to wait
36:   Set  $t_{\text{work}} \leftarrow \text{localTime}$ 
37:   Set  $\text{localTime} \leftarrow \text{localTime} + \text{shift}_i$  // Next slot in which staking will be performed is slot
     $\text{localTime} + 1$  according to the “new time”.
38: end if
39: end if

```

OUTPUT: The protocol outputs ok to its caller (but not to  $\mathcal{Z}$ ).

## C.12 Reading the Ledger State

Reading the ledger state is rather straightforward: first process all relevant information for this round and then extract the state.

### Protocol $\text{ReadState}(k, \mathcal{C}_{\text{loc}}, P, \text{sid}, R, f)$

- 1: If any of  $\text{isInit}$  or  $\text{isSync}$  is false output the empty state ( $\text{READ}, \text{sid}, \varepsilon$ ) (to  $\mathcal{Z}$ ). Otherwise, do the following:
- 2: Call  $\text{FetchInformation}(k, P)$  and denote the output by  $(\mathcal{C}_1, \dots, \mathcal{C}_M), (\text{tx}_1, \dots, \text{tx}_k)$ .
- 3: Set  $\text{buffer} \leftarrow \text{buffer} \parallel (\text{tx}_1, \dots, \text{tx}_k)$  and define  $\mathcal{N} \leftarrow \{\mathcal{C}_1, \dots, \mathcal{C}_M\}$
- 4: Call  $\text{UpdateTime}(P, R)$
- 5: Call  $\text{ProcessBeacons}(P, \text{sid})$
- 6: Let  $\mathcal{N}_0 := \{\mathcal{C} \in \mathcal{N} \cup \text{futureChains} \mid \forall B \in \mathcal{C} : \text{slotnum}(B) \leq \text{localTime}\}$
- 7: Let  $\mathcal{N}_1 := \{\mathcal{C} \in \mathcal{N} \mid \exists B \in \mathcal{C} : \text{slotnum}(B) > \text{localTime}\}$
- 8:  $\text{futureChains} \leftarrow (\text{futureChains} \setminus \mathcal{N}_0) \cup \mathcal{N}_1$
- 9:  $\text{fetchCompleted} \leftarrow \text{true}$
- 10: Call  $\text{SelectChain}(P, \text{sid}, \mathcal{C}_{\text{loc}}, \mathcal{N}_0, k, s, R, f)$  to update  $\mathcal{C}_{\text{loc}}$
- 11: Extract the state  $\text{st}$  from the current local chain  $\mathcal{C}_{\text{loc}}$ .
- 12: Output ( $\text{READ}, \text{sid}, \text{st}^{\lceil k}$ ) (to  $\mathcal{Z}$ ). //  $\text{st}^{\lceil k}$  denotes the prefix of  $\text{st}$  with the last  $k$  state blocks chopped off



### C.13 Simulate Clock Adjustments

Parties that have been absent only for a limited time (i.e., stalled for a limited time) can bootstrap very easily to the actual reliable state and time. Note that after doing this, they might be not yet alert for a small amount of time until they can start to work again, since they might have evolved their keys too far.

#### Protocol SimulateClockAdjustments( $P, R, k, f, s$ )

```

1: simulatedTime  $\leftarrow$  lastTimeAlert
2: for localTime  $-$  lastTimeAlert iterations do
3:   Let  $\mathcal{N}_0$  be the subsequence of futureChains s.t.  $\mathcal{C} \in \mathcal{N}_0 \Leftrightarrow \forall B \in \mathcal{C} : \text{slotnum}(B) \leq \text{simulatedTime}$ 
4:   Remove each  $\mathcal{C} \in \mathcal{N}_0$  from futureChains.
5:   Emulate SelectChain( $\mathcal{C}_{\text{loc}}, \mathcal{N}_0, k, s, R, f$ ) with simulated time simulatedTime (instead of localTime)
6:   - Update  $\mathcal{C}_{\text{loc}}$ 
7:   if simulatedTime mod  $R = 0$  then
8:     Emulate SyncProc( $P, R$ ) on simulated time simulatedTime (instead of localTime)
9:     - Execute Lines 1 to 13 to compute the shift  $\text{shift}_{\text{ep}}$  and to adjust already recorded arrival times.
10:    - Set simulatedTime  $\leftarrow$  simulatedTime +  $\text{shift}_{\text{ep}}$ 
11:   end if
12:   simulatedTime  $\leftarrow$  simulatedTime + 1
13: end for
14: Evolve the KES signing key by sending (USign, sid,  $P, 0, \text{localTime}$ ) to  $\mathcal{F}_{\text{KES}}$ 
15: Set  $t_{\text{work}} \leftarrow \text{localTime}$ 
16: Set localTime  $\leftarrow$  simulatedTime
OUTPUT: The protocol outputs ok to its caller (but not to  $\mathcal{Z}$ ).

```

### C.14 The round finish procedure

Once a party is done its actions in a round it has to advance the synchronous computation by sending the indication to  $\mathcal{G}_{\text{IMPERFLCLOCK}}$ . Since the functionality is shared, an update request in the ideal world will be relayed which implies that the protocol cannot simply ignore this input in the real world either. However, the update-request by the environment might not be well aligned with the round actions, so the protocol merely remembers that such an update has been received. At the end of its functions it then executes **FinishRound** which enforces that the protocol only sends the clock-update once (1) the round operations are concluded and (2) the environment has given the command to advance the round. (Note that in the ideal world, it is the ledger functionality which is registered with  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  and enforces the same principal time-evolving behavior as in the real world.)

#### Protocol FinishRound( $P$ )

```

1: while A (CLOCK-UPDATE,  $\text{sid}_C$ ) has not been received during the current round do
   Give up activation (set the anchor here)
end while
2: Send (CLOCK-UPDATE,  $\text{sid}_C$ ) to  $\mathcal{G}_{\text{IMPERFLCLOCK}}$ . // Party will lose its activation here.

```

### C.15 The Joining Procedure

One of the main procedures of Ouroboros Chronos concludes this description. The joining procedure will make any party that joins the system getting synchronized with the blockchain and to derive a local time-stamp that is in a small interval around the current alert parties time-stamps. The parameters of the procedure, along with their default values, are summarized in Table 1.

**Protocol JoinProc(P, sid, R, k, f, s, t<sub>off</sub>, t<sub>stable</sub>, t<sub>minSync</sub>)**

```

1: Call UpdateTime(P, R, f) // Align with newest round
2: if localTime > 1 then // Set back to local round 1
3:   Set localTime ← 1
4:   Set ep ← ⌈localTime/R⌉, and s1 ← localTime.
5:   fetchCompleted ← false, futureChains, buffer ← ∅, TimestampSB ← empty array.
6: end if
7: // Phase B
8: while localTime ≤ toff do
9:   if fetchCompleted = false then
10:    Call FetchInformation(k, P) and denote fetched chains by  $\mathcal{N} := (\mathcal{C}_1, \dots, \mathcal{C}_M)$ 
11:    Call SelectChain( $\mathcal{C}_{loc}, \mathcal{N}, k, s, R, f$ ) to update  $\mathcal{C}_{loc}$  // Since isSync = false, all chains are considered
12:    fetchCompleted ← true
13:    FinishRound(P) // Mark round actions as finished. Resume below upon next activation
14:   end if
15:   Call UpdateTime(P, R, f) to update localTime, ep, and s1 // fetchCompleted will reset.
16: end while
17: // Phases C
18: while localTime ≤ toff + tminSync + tstable do
19:   if fetchCompleted = false then
20:    Call FetchInformation(k, P) and denote the output by  $(\mathcal{C}_1, \dots, \mathcal{C}_M), (\mathbf{tx}_1, \dots, \mathbf{tx}_k)$ .
21:    Set buffer ← buffer ||  $(\mathbf{tx}_1, \dots, \mathbf{tx}_k)$  and define futureChains ← futureChains ||  $(\mathcal{C}_1, \dots, \mathcal{C}_M)$ 
22:    Call ProcessBeacons to collect new beacons in this round. // All arrival times are temporary
23:    Call SelectChain( $\mathcal{C}_{loc}, \text{futureChains}, k, s, R, f$ ) to update  $\mathcal{C}_{loc}$ 
24:    fetchCompleted ← true
25:    FinishRound(P) // Mark round actions as finished. Resume below upon next activation
26:   end if
27:   Call UpdateTime(P, R, f) to update localTime, ep, and s1 // fetchCompleted will reset.
28: end while
29: // Phase D
30: Define the function  $I_{\text{sync}}(j) : j \mapsto I_j := [(j-1)R+1, \dots, (j-1)R+2R/3]$ .
31: syncBuffervalid ← {SB' ∈ syncBuffer | ValidSB(P, sid, SB',  $\mathcal{C}_{loc}, f, R$ ) = true}
32: Initialize  $i := 0$ . Now set  $i$  to be the minimum positive integer such that
     $\forall \text{SB} \in \mathcal{C}_{loc}[I_{\text{sync}}(i)] : \text{SB} \in \text{syncBuffer}_{\text{valid}} \wedge \text{Timestamp}(\text{SB}) > t_{\text{off}} + t_{\text{pre}}$  (if no interval exists,  $i$  is unchanged).
33: if  $i \geq 1$  then
34:   for at most  $((t_{\text{stable}} + t_{\text{minSync}}) \text{div } R)$  iterations do
35:      $\mathcal{S}_i \leftarrow \{\text{SB} \mid \exists B \in \mathcal{C}_{loc}[I_{\text{sync}}(i)] : \text{SB} \in B \wedge \text{slotnum}(\text{SB}) \in \{(i-1)R+1, \dots, (i-1)R+R/6\}\}$ 
36:     for each  $\text{SB} = (s1, P, y, \pi) \in \mathcal{S}_i$  do
37:        $Q_{\text{SB}} \leftarrow \{\text{SB}' = (s1', P', \cdot, \cdot) \in \text{syncBuffer}_{\text{valid}} \mid P' = P \wedge s1' = s1\}$ 
38:       if  $Q_{\text{SB}} \neq \emptyset$  then
39:          $\min_{\text{SB}} \leftarrow \min\{\text{Timestamp}(\text{SB}') \mid \text{SB}' \in Q_{\text{SB}}\}$ 
40:          $\text{Timestamp}_{\text{SB}}(\text{SB}) \leftarrow (\min_{\text{SB}}, \text{final})$ 
41:          $\text{recom}(\text{SB}) \leftarrow \text{slotnum}(\text{SB}) - \text{Timestamp}(\text{SB})$ 
42:       else
43:          $S \leftarrow \mathcal{S} \setminus \{\text{SB}\}$  // Negligible probability event in execution.
44:       end if
45:     end for
46:      $\text{shift}_i \leftarrow \text{med}\{\text{recom}(\text{SB}) \mid \text{SB} \in \mathcal{S}_i\}$ 
47:     for each SB with  $\text{Timestamp}_{\text{SB}}(\text{SB}) = (a, \text{temp})$  do
48:        $\text{Timestamp}_{\text{SB}}(\text{SB}) \leftarrow (a + \text{shift}_i, \text{temp})$ 
49:     end for
50:     Set localTime ← localTime + shifti; EpochUpdate(i) ← Done
51:     Break if localTime ≤ (i+1)R. Otherwise, set  $i \leftarrow i+1$  and continue iteration.
52:   end for
53: isSync ← true; Run SelectChain( $\mathcal{C}_{loc}, \text{futureChains}, k, s, R, f$ ) to update  $\mathcal{C}_{loc}$ ;  $t_{\text{work}} \leftarrow \text{localTime} - 1$ 
54: If localTime ≤  $i \cdot R$  then set  $t_{\text{work}} \leftarrow i \cdot R$ . // Wait if shifted back before sync slot of  $i$ .
55: Evolve the KES signing key by sending (USign, sid, P, 0,  $t_{\text{work}}$ ) to  $\mathcal{F}_{\text{KES}}$ 
56: for each beacon  $\text{SB} \in \text{syncBuffer}_{\text{valid}}$  with  $\text{slotnum}(\text{SB}) \leq (i+1)R$  do
57:   Parse  $\text{Timestamp}_{\text{SB}}(\text{SB})$  as  $(a, \text{temp})$ . Define  $\text{Timestamp}_{\text{SB}}(\text{SB}) \leftarrow (a, \text{final})$ 
58: end for
59: end if

```

OUTPUT: The protocol outputs ok to its caller (but not to  $\mathcal{Z}$ ).

**Fig. 2.** The joining procedure.

Parameter	Default	Phase
$t_{\text{off}}$	$R/3$	B
$t_{\text{minSync}}$	$2R$	C
$t_{\text{stable}}$	$R$	C
$t_{\text{pre}}$	$3R/4$	D

**Table 1.** Parameters of the joining procedure and phases in which they play a role.

## D The Extended Ledger Functionality

As in [4], we prove composable security of this proof-of-stake protocol by showing that it realizes a ledger functionality that additionally exports additional time-stamps. This is important to show on what guarantees a higher-level protocol can rely as it requires to abstract the time advancement of the protocol in a way that is less complex than the real world. Compared to [4], we have thus two major differences:

1. The ledger uses  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  to maintain its baseline reference time (time since creation of the functionality).
2. We need to incorporate the detailed achieved guarantees for the exported time-stamps and provide a useful behavior for alert parties to be used by external protocols.

**Overview.** The ledger functionality introduced in [4] builds upon the general functionality [5]. In a nutshell, it maintains a central and unique ledger state denoted by `state`. How fast this state grows and which transactions it includes are part of the ledger policy—and therefore fully adjustable—as this depends heavily on the protocol achieving it. In any case, each registered party can request to see the state, and is guaranteed to receive a sufficiently long prefix of it; the size of each party’s view of the state is captured by (monotonically) increasing pointers that define which part of the state each party can read; the adversary has a limited control on these pointers. The dynamics of this can be reflected with a sliding window over the sequence of state blocks, with width `windowSize` and starting at the head of the state; each party’s pointer points to a location within this window. The adversary can choose the position of the pointers within this sliding window. As is common in UC, parties advance the ledger when they are instructed to (activated with specific maintain-ledger input by their environment  $\mathcal{Z}$ ). The ledger uses these queries along with the function `predict-time( $\cdot$ )` to ensure that the ideal world execution advances with the same pace (relatively to the pacemaker)—in this work the *baseline speed captured by the nominal time derived from  $\mathcal{G}_{\text{IMPERFLCLOCK}}$* —as the protocol.

Any party can input a transaction to the ledger; upon reception transactions are validated using a predicate `Validate` and, if found valid, are added to a buffer. Each new block of the state consists of transactions which were previously accepted to the buffer. To give protocols syntactic freedom to structure their state blocks, a vector of transactions, say  $\mathbf{N}_i$ , is mapped to the  $i$ th state block via function `Blockify( $\mathbf{N}_i$ )`. `Validate` and `Blockify` are two of the ledger’s parametrization algorithms. The third algorithm is the predicate `predict-time` that is instantiated by the predictable time-behavior predicate of the protocol under consideration (in this case Ouroboros Chronos). We note that that Ouroboros Chronos has a predictable time-advancement pattern due to the way the UC protocol is designed and it is always clear when honest parties call `FinishRound` in the protocol.

**Party sets maintained by the ledger.** In accordance with the classification of parties in Section 4, the ledger divides the registered honest into two different basic categories called *synchronized* and *desynchronized*. Synchronized parties are the parties that enjoy full security guarantees. Formally, a party that is considered synchronized by the ledger and which is connected to all shared setups is what we usually called *alert* in the dynamic availability setting. A party is considered synchronized if it has been continuously connected to all its resources for a sufficiently long interval and has maintained connectivity to these resources until the current time. Formally, here, “sufficiently long” refers to `Delay`-many rounds, where `Delay` is a parameter of the ledger and directly relates to the real-world time duration between joining and becoming synchronized. In the general formulation, de-synchronized parties receive significantly weaker guarantees.

**State-extend policy.** A defining part of the behavior of the ledger is the (parameterizable) procedure which defines when/how to extend `state` and what the constraints are for an adversary. The state-extend policy of the ledger in this work is almost identical to the ledger in [4] except that we establish a slightly better bound for transaction liveness guarantees. In nutshell, the basic mode of operation of `ExtendPolicy` is that it takes as an input a proposal from the adversary for extending the state, and can decide to follow this proposal if it satisfies its policy; if it does not, `ExtendPolicy` can ignore the proposal (and enforce a default extension). It will enforce minimal chain growth, a certain fraction of “good blocks,” and transaction liveness guarantees for old and (still) valid transactions.

Given this similarity, we give a concise summary of the ledger parameters used in this work in Table M.2. The ledger functionality and its `ExtendPolicy` are further provided in the appendix for the sake of self-containment.

**The Export-Time Extension.** We introduce the export time extension to  $\mathcal{G}_{\text{LEDGER}}$ . first, we represent a time-stamps `timeP` associated to party  $P$  as a pair  $(e, t)$ , where  $t$  is the actual time stamp, and  $e$  refers to what we call a epoch.<sup>20</sup> An alert party’s time  $t$  in  $(e, t)$  is guaranteed to increase during an epoch with every tick of the reference speed. Once  $t$  hits an epoch boundary, defined as multiples of an epoch length parameter  $R_L$ <sup>21</sup>, the epoch value increases as well. Clearly, this would a perfect, monotonically increasing, two-dimensional time-stamp. We have to weaken this guarantee by allowing to the adversary to apply a limited shift whenever a party is at an epoch boundary (parameters `shiftLB`, `shiftUB`). Furthermore the ledger enforces that any two alert parties with respective time-stamps  $(e, t)$  and  $(e', t')$ , satisfy the constraints  $|t - t'| \leq \text{timeSlack}_{\text{total}}$  and  $|t - t'| \leq \text{timeSlack}_{\text{ep}}$  if  $e = e'$ , and  $|e - e'| \leq 1$  for the respective ledger parameters `timeSlackep`, `timeSlacktotal` that define the maximally allowed skewness of parties. Note that we give the possibility than Within an epoch the slack could be potentially different (i.e., much better) than across epochs.

We give an overview of the parameters in Table M.2 and provide the functionality in Section D.1.

## D.1 The Functionality

For completeness, we describe here the ledger functionality in the  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  world together with the export-time extension. A couple of technicalities might be mentioned here regarding the novel aspects for time-management: as explained in the main body, the timestamps are two-dimensional of the form  $(e, t)$ , where  $e$  is an epoch number and  $t$  is the actual timestamp. Initially, the epoch is  $e = -1$  which stands for “prior” to be active. Timestamps advance in a monotone fashion controlled by the functionality. The only irregularities are introduced, by the adversary, once the epoch increases from  $e$  to  $e + 1$  for a party. Then the functionality allows the adversary shifting the reported timestamp of that party by a limited amount. The influence is captured by the above explained clock parameters. The first switch the adversary can do (when  $e = -1$  and  $t = 0$ ) corresponds to the initial offset. We highlight the differences to the PoS ledger from [4] in blue to make the quite complex definition more modular.

### Functionality $\mathcal{G}_{\text{LEDGER}}$ with Export-Clock Extension

**General:** The functionality is parameterized by four algorithms, `Validate`, `ExtendPolicy`, `Blockify`, and `predict-time`, along with three parameters: `windowSize`, `Delay`  $\in \mathbb{N}$ , and  $\mathcal{S}_{\text{initStake}} := \{(U_1, s_1), \dots, (U_n, s_n)\}$ . The functionality manages variables `state`, `NxtBC`, `buffer`,  $\tau_L$ , and  $\tau_{\text{state}}$ , as described above. The variables are initialized as follows: `state` :=  $\tau_{\text{state}}$  := `NxtBC` :=  $\varepsilon$ , `buffer` :=  $\emptyset$ ,  $\tau_L = 0$ . For each party  $P \in \mathcal{P}$  the functionality maintains a pointer `pti` (initially set to 1) and a current-state view `statep` :=  $\varepsilon$  (initially set to empty). The functionality also keeps track of the timed honest-input sequence in a vector  $\mathcal{I}_H^T$  (initially  $\mathcal{I}_H^T := \varepsilon$ ).

**Party Management:** The functionality maintains the set of registered parties  $\mathcal{P}$ , the (sub-)set of honest parties  $\mathcal{H} \subseteq \mathcal{P}$ , and the (sub-set) of de-synchronized honest parties  $\mathcal{P}_{DS} \subset \mathcal{H}$  (as discussed below). The sets  $\mathcal{P}$ ,  $\mathcal{H}$ ,  $\mathcal{P}_{DS}$  are

<sup>20</sup> In the real world,  $e$  would be the number of adjustments a party has made to its local clock.

<sup>21</sup> In this work, this is always exactly the length of normal epoch, i.e.,  $R_L = R$ .

all initially set to  $\emptyset$ . When a (currently unregistered) honest party is registered at the ledger, *if it is registered with the clock-tick functionality and the global RO already*, then it is added to the party sets  $\mathcal{H}$  and  $\mathcal{P}$  and the current time of registration is also recorded; if the current time is  $\tau_L > 0$ , it is also added to  $\mathcal{P}_{DS}$ . Similarly, when a party is deregistered, it is removed from both  $\mathcal{P}$  (and therefore also from  $\mathcal{P}_{DS}$  or  $\mathcal{H}$ ). The ledger maintains the invariant that it is registered (as a functionality) to  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  in round  $\tau_L = 0$  or whenever  $\mathcal{H} \neq \emptyset$  in which case it updates the clock at the speed of the slowest party in  $\mathcal{H}$ .

**Time management:** When activated with the first registration command, the ledger prepends to its actions above the following steps: it asks for the clock tick to  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  and stores the variable internally as `lastTick` and sets  $\tau_L := 0$ .

It further activates the adversary with restricting query (`Respond`, (`start`, `sid`)) and resumes with normal registration after receiving the immediate acknowledgment from the adversary.

**Handling initial stakeholders:** The ledger blocks the advancement of the baseline time of the clock until each initial stakeholder, i.e.,  $P \in \mathcal{S}_{\text{initStake}}$  has registered. Additionally, the adversary can assign a unique reference string `refP` once to each registered party  $P$ .

**Extension Parameters:** The extension is parameterized by `shiftLB`, `shiftUB`, `timeSlackep`, `timeSlacktotal`, and the epoch length  $R_L$ .

**Extension Variables:** The extension introduces the new variables `timeP` for each registered party  $P \in \mathcal{P}$  (initial value  $(-1, 0)$ ).

**Upon receiving any input  $I$**  from any party or from the adversary, send (`CLOCK-GET`, `sidC`) to  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  and upon receiving response (`CLOCK-GET`, `sidC`, `tick`) set  $d = 1$  if `lastTick`  $\neq$  `tick` and 0 otherwise. Set  $\tau_L := \tau_L + d$ . Do the following if  $\tau_L > 0$  (otherwise, ignore input):

1. Updating synchronized/desynchronized party set:
  - (a) Let  $\widehat{\mathcal{P}} \subseteq \mathcal{P}_{DS}$  denote the set of desynchronized honest parties that have been registered (continuously) to the ledger,  $\mathcal{G}_{\text{IMPERFLCLOCK}}$ , and the GRO since time  $\tau' < \tau_L - \text{Delay}$ .
  - (b) For each party  $P \in \widehat{\mathcal{P}}$  with `timeP` =  $(e, t)$ , ensure valid range of timestamps: verify that  $t \leq \tau_L + \Delta_{\text{clock}}$  and that for any party  $P' \in \mathcal{H} \setminus \mathcal{P}_{DS}$  with `timeP'` =  $(e', t')$ , it holds that  $|t - t'| \leq \text{timeSlack}_{\text{total}}$ , if  $e = e'$  it also holds that  $|t - t'| \leq \text{timeSlack}_{\text{ep}}$  and that  $|e - e'| \leq 1$ . If a comparison fails, set `timeP` to that value `timeP'`.
  - (c) Set  $\mathcal{P}_{DS} := \mathcal{P}_{DS} \setminus \widehat{\mathcal{P}}$ .
  - (d) For any synchronized party  $P \in \mathcal{H} \setminus \mathcal{P}_{DS}$ , if  $P$  is not registered to the clock, then consider it desynchronized, i.e., set  $\mathcal{P}_{DS} \cup \{P\}$ .
2. If  $I$  was received from an honest party  $P \in \mathcal{P}$ :
  - (a) Set  $\mathcal{I}_H^T := \mathcal{I}_H^T \parallel (I, P, \tau_L)$ ;
  - (b) Compute  $N = (N_1, \dots, N_\ell) := \text{ExtendPolicy}(\mathcal{I}_H^T, \text{state}, \text{NxtBC}, \text{buffer}, \tau_{\text{state}})$  and if  $N \neq \varepsilon$  set `state` := `state` || `Blockify`( $N_1$ ) || ... || `Blockify`( $N_\ell$ ) and  $\tau_{\text{state}} := \tau_{\text{state}} \parallel \tau_L^\ell$ , where  $\tau_L^\ell = \tau_L \parallel \dots \parallel \tau_L$ .
  - (c) For each `BTX`  $\in$  `buffer`: if `Validate`(`BTX`, `state`, `buffer`) = 0 then delete `BTX` from `buffer`. Also, reset `NxtBC` :=  $\varepsilon$ .
  - (d) If there exists  $U_j \in \mathcal{H} \setminus \mathcal{P}_{DS}$  such that  $|\text{state}| - \text{pt}_j > \text{windowSize}$  or  $\text{pt}_j < |\text{state}_j|$ , then set  $\text{pt}_k := |\text{state}|$  for all  $U_k \in \mathcal{H} \setminus \mathcal{P}_{DS}$ .
3. Increase the party time stamps in a new round:
  - (a) For all  $P \in \mathcal{H} \setminus \mathcal{P}_{DS}$  do: parse `timeP` as  $(e, t)$  and set `timeP`  $\leftarrow$   $(e, t + 1)$  for all the parties that have advanced locally according to  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  since the last activation. If  $P$  is stalled and  $t + 1 \text{ div } R_L > e$  then  $\mathcal{P}_{DS} \cup \{P\}$ .
4. If the calling party  $P$  is *stalled* or *time-unaware* (according to the defined party classification), then no further actions are taken. Otherwise, depending on the above input  $I$  and its sender's ID,  $\mathcal{G}_{\text{LEDGER}}$  executes the corresponding code from the following list:

- *Submitting a transaction:*  
If  $I = (\text{SUBMIT}, \text{sid}, \text{tx})$  and is received from a party  $P \in \mathcal{P}$  or from  $\mathcal{A}$  (on behalf of a corrupted party  $P$ ) do the following
  - Choose a unique transaction ID  $\text{txid}$  and set  $\text{BTX} := (\text{tx}, \text{txid}, \tau_L, P)$
  - If  $\text{Validate}(\text{BTX}, \text{state}, \text{buffer}) = 1$ , then  $\text{buffer} := \text{buffer} \cup \{\text{BTX}\}$ .
  - Send  $(\text{SUBMIT}, \text{BTX})$  to  $\mathcal{A}$ .
- *Reading the state:*  
If  $I = (\text{READ}, \text{sid})$  is received from a party  $P \in \mathcal{H} \setminus \mathcal{P}_{DS}$  then immediately return  $(\text{READ}, \text{sid}, \varepsilon)$  if  $\text{time}_P = (e, t)$  with  $e \leq 0 \wedge t < 0$ . Else, set  $\text{state}_P := \text{state}|_{\min\{\text{pt}_P, |\text{state}|\}}$  and return  $(\text{READ}, \text{sid}, \text{state}_P)$  to the requester. If the requester is  $\mathcal{A}$  then send  $(\text{state}, \text{buffer}, \mathcal{I}_H^T)$  to  $\mathcal{A}$ .
- *Maintaining the ledger state:*  
If  $I = (\text{MAINTAIN-LEDGER}, \text{sid}, \text{minerID})$  is received by an honest party  $P \in \mathcal{P}$  and (after updating  $\mathcal{I}_H^T$  as above)  $\text{predict-time}(\mathcal{I}_H^T, \text{aux}) = \hat{\tau} > \tau_L$  then send  $(\text{CLOCK-UPDATE}, \text{sid}_C)$  to  $\mathcal{G}_{\text{IMPERFLCLOCK}}$ . Here we provide  $\text{aux} = (\text{state}, (P_1, \text{ref}_{P_1}, \text{time}_{P_1}), \dots, (P_{|\mathcal{P}|}, \text{ref}_{P_{|\mathcal{P}|}}, \text{time}_{P_{|\mathcal{P}|}}))$ . Else, send  $I$  to  $\mathcal{A}$ .  
Before losing the activation, if party  $P \in \mathcal{H} \setminus \mathcal{P}_{DS}$  has only locally completed its round actions, update the clock in the name of this party. If  $\text{time}_P = (e, t)$  for  $t > 0$  and  $t \text{ div } R_L > e$  then set  $\text{time}_P \leftarrow (e + 1, t)$ .
- *The adversary proposing the next block:*  
If  $I = (\text{NEXT-BLOCK}, \text{hFlag}, (\text{txid}_1, \dots, \text{txid}_\ell))$  is sent from the adversary, update  $\text{NxtBC}$  as follows:
  - Set  $\text{listOfTxid} \leftarrow \varepsilon$
  - For  $i = 1, \dots, \ell$  do: if there exists  $\text{BTX} := (x, \text{txid}, \text{minerID}, \tau_L, U_j) \in \text{buffer}$  with ID  $\text{txid} = \text{txid}_i$  then set  $\text{listOfTxid} := \text{listOfTxid} \parallel \text{txid}_i$ .
  - Finally, set  $\text{NxtBC} := \text{NxtBC} \parallel (\text{hFlag}, \text{listOfTxid})$  and output  $(\text{NEXT-BLOCK}, \text{ok})$  to  $\mathcal{A}$ .
- *The adversary setting state-slackness:*  
If  $I = (\text{SET-SLACK}, (U_{i_1}, \widehat{\text{pt}}_{i_1}), \dots, (U_{i_\ell}, \widehat{\text{pt}}_{i_\ell}))$ , with  $\{P_{i_1}, \dots, P_{i_\ell}\} \subseteq \mathcal{H} \setminus \mathcal{P}_{DS}$  is received from the adversary  $\mathcal{A}$  do the following:
  - If for all  $j \in [\ell] : |\text{state}| - \widehat{\text{pt}}_{i_j} \leq \text{windowSize}$  and  $\widehat{\text{pt}}_{i_j} \geq |\text{state}_{i_j}|$ , set  $\text{pt}_{i_1} := \widehat{\text{pt}}_{i_1}$  for every  $j \in [\ell]$  and return  $(\text{SET-SLACK}, \text{ok})$  to  $\mathcal{A}$ .
  - Otherwise set  $\text{pt}_{i_j} := |\text{state}|$  for all  $j \in [\ell]$ .
- *The adversary setting the state for desynchronized parties:*  
If  $I = (\text{DESYNC-STATE}, (U_{i_1}, \text{state}'_{i_1}), \dots, (U_{i_\ell}, \text{state}'_{i_\ell}))$ , with  $\{U_{i_1}, \dots, U_{i_\ell}\} \subseteq \mathcal{P}_{DS}$  is received from the adversary  $\mathcal{A}$ , set  $\text{state}_{i_j} := \text{state}'_{i_j}$  for each  $j \in [\ell]$  and return  $(\text{DESYNC-STATE}, \text{ok})$  to  $\mathcal{A}$ .
- *Reading the time:*  
If  $I = (\text{EXPORT-TIME}, \text{sid})$  is received from a party  $P \in \mathcal{H} \setminus \mathcal{P}_{DS}$  then parse  $\text{time}_P$  as  $(e, t)$ . If  $e \geq 0$  and  $t \geq 0$  then return  $(\text{EXPORT-TIME}, \text{sid}, \text{time}_P)$  to the requester. Otherwise, return  $(\text{EXPORT-TIME}, \text{sid}, \perp)$ .
- *The adversary setting the shift on epoch boundaries :*  
If  $I = (\text{APPLY-SHIFT}, \text{sid}, (P, s))$  is received from the adversary  $\mathcal{A}$  and  $P \in \mathcal{H} \setminus \mathcal{P}_{DS}$  then do the following:
  - Parse  $\text{time}_P$  as  $(e, t)$ . If  $t \text{ mod } R_L \neq 0$  return to the adversary.
  - Verify that  $\text{shiftLB} \leq s \leq \text{shiftUB}$ . If the check fails, return to the adversary.
  - Set  $t' \leftarrow t + s$  and verify correct range of timestamps:  
Check that  $t' \leq \tau_L$  and that for each party  $P' \in \mathcal{H} \setminus \mathcal{P}_{DS}$  with  $\text{time}_{\text{time}'_P} = (e'', t'')$ , it holds that  $|t' - t''| \leq \text{timeSlack}_{\text{total}}$ ,  $|e' - e''| \leq 1$ , and if  $e' = e''$  verify that  $|t' - t''| \leq \text{timeSlack}_{\text{ep}}$ .
  - If all checks succeeds, set  $\text{time}_P \leftarrow (e, t')$ . Otherwise, set  $\text{time}_P \leftarrow \text{time}_{P'}$  where  $P'$  is the lexicographically smallest identity of  $\mathcal{H} \setminus \mathcal{P}_{DS}$ . Return activation to the adversary.
- *The adversary setting the timestamp for desynchronized parties:*  
If  $I = (\text{SET-TIME}, \text{sid}, P, \text{time})$  is received from the adversary  $\mathcal{A}$  do the following: if  $P \in \mathcal{P}_{DS}$  then set  $\text{time}_P \leftarrow \text{time}$

## D.2 Extend Policy

For completeness, we state here the extend policy of [4] and mark in blue the minor modification. Note that the default block mechanism DEFAULTEXTENSION is identical to [4] and thus omitted.

### Algorithm ExtendPolicy for $\mathcal{G}_{\text{LEDGER}}$

```

function EXTENDPOLICY( $\mathcal{I}_H^T$ , state, NxtBC, buffer,  $\tau_{\text{state}}$ )
  // First, create a default honest client block as alternative:
   $N_{\text{df}} \leftarrow \text{DEFAULTEXTENSION}(\mathcal{I}_H^T, \text{state}, \text{NxtBC}, \text{buffer}, \tau_{\text{state}})$  // Extension if adversary violates policy.
  Let  $\tau_L$  be current ledger time (computed from  $\mathcal{I}_H^T$ )
  // The function must not have side-effects: Only modify copies of relevant values.
  Create local copies of the values buffer, state, and  $\tau_{\text{state}}$ .
  // Now, parse the proposed block by the adversary
  Parse NxtBC as a vector  $((\text{hFlag}_1, \text{NxtBC}_1), \dots, (\text{hFlag}_n, \text{NxtBC}_n))$ 
   $N \leftarrow \varepsilon$  // Initialize Result
  // Determine the time of the state block which is windowSize blocks behind the head of the state
  if  $|\text{state}| \geq \text{windowSize}$  then
    Set  $\tau_{\text{low}} \leftarrow \tau_{\text{state}}[|\text{state}| - \text{windowSize} + 1]$ 
  else
    Set  $\tau_{\text{low}} \leftarrow 1$ 
  end if
  oldValidTxMissing  $\leftarrow$  false // Flag to keep track whether old enough, valid transactions are inserted.
  for each list  $\text{NxtBC}_i$  of transaction IDs do
    // Compute the next state block
    // Verify validity of  $\text{NxtBC}_i$  and compute content
    Use the txid contained in  $\text{NxtBC}_i$  to determine the list of transactions
    Let  $\text{tx} = (\text{tx}_1, \dots, \text{tx}_{|\text{NxtBC}_i|})$  denote the transactions of  $\text{NxtBC}_i$ 
    if  $\text{tx}_1$  is not a coin-base transaction then
      return  $N_{\text{df}}$ 
    else
       $N_i \leftarrow \text{tx}_1$ 
      for  $j = 2$  to  $|\text{NxtBC}_i|$  do
        Set  $\text{st}_i \leftarrow \text{blockify}_{\mathbb{B}}(N_i)$ 
        if  $\text{ValidTx}_{\mathbb{B}}(\text{tx}_j, \text{state} || \text{st}_i) = 0$  then
          return  $N_{\text{df}}$ 
        end if
         $N_i \leftarrow N_i || \text{tx}_j$ 
      end for
      Set  $\text{st}_i \leftarrow \text{blockify}_{\mathbb{B}}(N_i)$ 
    end if
    // Test that all old valid transaction are included
    if the proposal is declared to be an honest block, i.e.,  $\text{hFlag}_i = 1$  then
      for each  $\text{BTX} = (\text{tx}, \text{txid}, \tau', \text{P}) \in \text{buffer}$  of an honest party P with time  $\tau' < \tau_{\text{low}} - \text{Delay}_{\text{tx}}$  do
        if  $\text{ValidTx}_{\mathbb{B}}(\text{tx}, \text{state} || \text{st}_i) = 1$  but  $\text{tx} \notin N_i$  then
          oldValidTxMissing  $\leftarrow$  true
        end if
      end for
    end if
     $N \leftarrow N || N_i$ 
     $\text{state} \leftarrow \text{state} || \text{st}_i$ 
     $\tau_{\text{state}} \leftarrow \tau_{\text{state}} || \tau_L$ 
    // Must not proceed with too many adversarial blocks
     $i \leftarrow \max\{\{\text{windowSize}\} \cup \{k \mid \text{st}_k \in \text{state} \wedge \text{proposal of } \text{st}_k \text{ had } \text{hFlag} = 1\}\}$  // Determine most
    // recent honestly-generated block in the interval behind the head.
  
```

```

if |state| - i ≥ advBlckswindow then
  return Ndf
end if
end if
// Update τlow: the time of the state block which is windowSize blocks behind the head of the
// current, potentially already extended state
if |state| ≥ windowSize then
  Set τlow ← τstate[|state| - windowSize + 1]
else
  Set τlow ← 1
end if
end for
// Final checks (if policy is violated, it is enforced by the ledger):
// Must not proceed too slow or with missing transaction.
if τlow > 0 and τL - τlow > maxTimewindow then // A sequence of blocks cannot take too much time.
  return Ndf
else if τlow = 0 and τL - τlow > 2 · maxTimewindow then // Bootstrapping cannot take too much time.
  return Ndf
else if oldValidTxMissing then // If not all old enough, valid transactions have been included.
  return Ndf
end if
return N
end function

```

## E On Using the Exported Clock

The goal of this section is to give an overview on how external protocols could make use of the timing service provided by Ouroboros Chronos. For simplicity, this section assume that all parties run at equal speed, i.e.,  $\Delta_{\text{clock}} = 0$ . In this model, the nominal time has a simple meaning: it is the objective number of rounds passed in the execution.

### E.1 General Considerations

Cryptographic protocols can use the exported clock of Chronos and make use of the provided timestamps. It is instructive to see different cases depending on the parameters of the clock. For example, if  $\text{timeSlack}_{\text{ep}} = \text{timeSlack}_{\text{total}} = 0$ , and  $\text{shiftLB} = \text{shiftUB} = 0$ , then we have an equivalent formulation of the global clock of previous works. Each weakening of the parameters will result in a higher-level protocol to require specific reactions. This is, however, possible: for example, if shifts are to be expected but still  $\text{timeSlack}_{\text{ep}} = \text{timeSlack}_{\text{total}} = 0$ , then the parties will know that some strange behavior could happen around the epoch boundaries. However, the behavior is limited and predictable based on the clock parameters. For example, parties could stall their operations just before the epoch boundary switches and depending on the shift, resume their operations later at a specific time. Furthermore, by the limited shift and the guaranteed advancement the parties will proceed and, if the protocol uses explicit knowledge of  $\text{shiftLB}$  and  $\text{shiftUB}$ , liveness can be explicitly quantified. If parties can further be skewed, in addition to the above, the higher level protocol has to be resilient against small variations in the time-stamps. Again, the level of resilience required is clearly defined by parameters  $\text{timeSlack}_{\text{total}}$  and  $\text{timeSlack}_{\text{ep}}$ .

### E.2 Clock Properties in Optimistic Network Models

In the above cryptographic treatment, we made worst-case assumptions regarding the delivery times of the synchronization beacons and chains. One can ask the question how well Chronos adjusts to more optimistic network models, i.e., how the clock parameters get more precise if Chronos is executed in a less adversarial environment. We can infer some rough estimates based on our analysis by formally imposing restrictions on the delays occurring in the network and then reworking the derivation of the main clock properties along the



lines of Lemma 9. We exemplify this with two examples. Let us assume that we have initial coordination, i.e., that all parties obtain the genesis block in the same objective round. Let us further assume that all messages sent in round  $\ell$  are guaranteed to be delivered in round  $\ell + 1$ —or more generally in round  $\ell + m$ . Along the lines of Lemma 9, we observe that the timestamps reported by alert parties are very coordinated. In fact, if  $\tau$  denotes objective time (as recorded inside  $\mathcal{G}_{\text{IMPERFLCLOCK}}$ ), then each alert party reports time  $t = \tau - m \cdot e$  (where  $e$  is the timestamp’s epoch number) and thus the clock parameters of the ledger would in this case be  $\text{shiftLB} = \text{shiftUB} = -m$  and  $\text{timeSlack}_{\text{ep}} = 0$ ,  $\text{timeSlack}_{\text{total}} = m$ . Even more, in this case, any external protocol can recompute the objective time (as recorded by  $\mathcal{G}_{\text{IMPERFLCLOCK}}$ ) from a timestamp  $(e, t)$  of an alert party by the adjustment  $t + e \cdot m$  and thereby obtain a perfect “global clock” whenever this particular “network parameter”  $m$  is known.

More realistically, if we consider delays in an interval of the form  $[m - d, m + d] \subseteq [0, \Delta]$  then the arguments of Lemma 9 suggest the clock parameters  $\text{shiftLB} = -m - 3d$ ,  $\text{shiftUB} = -m + 3d$  and  $\text{timeSlack}_{\text{ep}} = 2d$ ,  $\text{timeSlack}_{\text{total}} = m + 3d$ . In this case, the above proposed time-adjustment  $t + e \cdot m$ , based on the characteristic value  $m$  of the network, would yield an approximation of the objective  $\tau$  in the order of the width  $2d$  of the interval (note that this naturally matches our security analysis when choosing  $m = d$  and  $2d = \Delta$ , i.e., when  $m$  is the center of  $[0, \Delta]$ ).

In the above considerations, the term  $2d$  refers to a logical width measured in number of rounds. Depending on the message sizes and the physical time duration of one round, if messages that a party needs to send during round  $\ell$  are always sent at the end of round  $\ell$ , then a logical width of  $2d \approx 0$  might not be unrealistic to achieve. Also, an accurate approximation of the “typical delay” parameter  $m$  seems in principle feasible. Hence, obtaining quite accurate approximations of objective time seems achievable by Ouroboros Chronos in networks that have a somewhat predictable behavior. This, however, depends heavily on the implementations and the network stack and a more accurate study is part of future work.

## F Single-Epoch Security with Static Registration and $\Delta$ -Bounded Skew

In this section we prove Theorem 1. First, in Section F.1, we introduce a simplified chain-selection rule that will make our protocol easier to analyze. In Section F.2 we draw the connection between a single-epoch execution of this simplified protocol and the formalism of characteristic strings and forks that we later employ. We then analyze the distribution of the characteristic strings induced by an execution of the simplified Ouroboros Chronos in Section F.3, and establish the implications of that for the properties CP, CG and CQ in Section F.4. Finally, in Section F.5, we replace the simplified chain-selection rule with the actual one, concluding the proof of Theorem 1.

### F.1 The Simplified Chain-Selection Rule `maxvalid-mc`

To make our analysis more modular, and take advantage of the combinatorial framework for analyzing common-prefix violations of longest-chain rule protocols developed gradually in [26,14,4], we first consider the protocol Ouroboros-Chronos with a simplified chain-selection rule `maxvalid-mc` (given in Fig. 3) instead of the actual rule `maxvalid-bg` given in Fig. C.8; we will denote this variant `Ouroboros-Chronosmc` for conciseness.

The rule `maxvalid-mc` differs in that it applies the longest-chain preference and refuses to revert more than  $k$  blocks under any circumstances (hence the “mc” identifier standing for “moving checkpoint”). This is in contrast to the more nuanced behavior of `maxvalid-bg` that compares the two chains forking more than  $k$  blocks ago for density close to the point where they fork (cf. Condition B in Fig. C.8). The latter rule allows for so-called *bootstrapping from genesis* [4] (hence “bg”) and so we adopt it for Ouroboros Chronos as well, the consequences for our analysis are discussed in Section F.5.

### F.2 From Executions to Forks

We recall the notion of a *characteristic string*, which we use to record, for each slot in a sequence of slots, whether any leader is elected for the slot and, if that is the case, whether this leader is unique and alert.

**Protocol maxvalid-mc**( $C_{\text{loc}}, C_1, \dots, C_\ell$ )

```

1: Set  $C_{\text{max}} \leftarrow C_{\text{loc}}$ .
2: for  $i = 1$  to  $\ell$  do
3:   if  $\text{IsValidChain}(C_i)$  then
  // Compare  $C_{\text{max}}$  to  $C_i$ 
4:     if ( $C_i$  forks from  $C_{\text{max}}$  at most  $k$  blocks) then
5:       if  $|C_i| > |C_{\text{max}}|$  then // Condition A
         Set  $C_{\text{max}} \leftarrow C_i$ .
       end if
     end if
  end if
  end for
6: return  $C_{\text{max}}$ .

```

**Fig. 3.** The simplified chain selection rule maxvalid-mc.

**Definition 6 (Characteristic string [26,14,4]).** Let  $S = \{\text{sl}_1, \dots, \text{sl}_R\}$  be a sequence of slots of length  $R$ ; consider an execution (with adversary  $\mathcal{A}$  and environment  $\mathcal{Z}$ ) of the protocol. For a slot  $\text{sl}_j$ , let  $P(j)$  denote the set of active parties assigned to be slot leaders for slot  $j$  by the protocol. We define the characteristic string  $w \in \{0, 1, \perp\}^R$  of  $S$  to be the random variable so that

$$w_j = \begin{cases} \perp & \text{if } P(j) = \emptyset, \\ 0 & \text{if } |P(j)| = 1 \text{ and the assigned party is alert,} \\ 1 & \text{otherwise.} \end{cases} \quad (7)$$

For such a characteristic string  $w \in \{0, 1, \perp\}^*$  we say that the index  $j$  is uniquely alert if  $w_j = 0$ , empty if  $w_j = \perp$ , and potentially active if  $w_j \in \{0, 1\}$ .

If the execution is fixed (i.e., the randomness of the execution is fixed), we use the notation  $w_{\mathcal{E}}$  to denote the fixed characteristic string resulting from that particular execution, where the subscript  $\mathcal{E}$  is used to indicate its difference to the random variable above.

We also recall the notion of a  $\Delta$ -fork, a tool developed to reason about the various blockchains that can be induced by an adversary in the  $\Delta$ -synchronous setting with a particular characteristic string.

**Definition 7 ( $\Delta$ -fork [14]).** Let  $w \in \{0, 1, \perp\}^k$  and  $\Delta$  be a non-negative integer. Let  $A = \{i \mid w_i \neq \perp\}$  denote the set of potentially active indices, and let  $H = \{i \mid w_i = 0\}$  denote the set of uniquely alert indices. A  $\Delta$ -fork for the string  $w$  is a rooted tree  $F = (V, E)$  with a labeling  $\ell : V \rightarrow \{0\} \cup A$  so that

- (i) the root  $r \in V$  is given the label  $\ell(r) = 0$ ;
- (ii) the labels along any (simple) path beginning at the root are strictly increasing;
- (iii) each uniquely alert index  $i \in H$  is the label of exactly one vertex of  $F$ ;
- (iv) the function  $\mathbf{d} : H \rightarrow \{1, \dots, k\}$ , defined so that  $\mathbf{d}(i)$  is the depth in  $F$  of the unique vertex  $v$  for which  $\ell(v) = i$ , satisfies the following  $\Delta$ -monotonicity property: if  $i, j \in H$  and  $i + \Delta < j$ , then  $\mathbf{d}(i) < \mathbf{d}(j)$ .

For convenience, we direct the edges of forks so that depth increases along each edge; then there is a unique directed path from the root to each vertex and, in light of (ii), labels along such a path are strictly increasing. As a matter of notation, we write  $F \vdash_{\Delta} w$  to indicate that  $F$  is a  $\Delta$ -fork for the string  $w$ . We typically refer to a  $\Delta$ -fork as simply a “fork”.

Note that both notions of a characteristic string and a fork can be directly ported to our setting without a global clock, interpreting the slot indices as *logical time*, in accordance with the rest of this paper (cf. Section 2). However, this change of the setting requires us to re-establish the connection between executions and forks from [14]. The relevant part of the outcome of an execution is captured in the notion of an *execution tree* which we first define, the transition from executions to forks is then stated in Lemma 2.

**Definition 8 (Execution tree [14]).** Consider an execution  $\mathcal{E}$  of the real-world experiment. The execution tree for this execution is a directed, rooted tree  $T_{\mathcal{E}} = (V, E)$  with a labeling  $\ell : V \rightarrow \mathbb{N}_0$  that is constructed during the execution as follows:

- (i) At the beginning,  $V = \{r\}$ ,  $E = \emptyset$  and  $\ell(r) = 0$ .
- (ii) Every chain  $C'$  that is input to `maxvalid-bg` as a part of  $\mathcal{N}$  or created as a new local chain in Step 20 of `StakingProcedure` of `Ouroboros-Chronos` run by any alert party is immediately processed block-by-block from the genesis block to  $\text{head}(C')$ . For every block  $B = (h, \text{st}, \text{sb}, \text{s1}, \text{crt}, \rho, \sigma)$  processed for the first time:
  - a new vertex  $v_B$  is added to  $V$ ;
  - a new edge  $(v_{B^-}, v_B)$  is added to  $E$  where  $B^-$  is the unique block such that  $H(B^-) = h$ ;
  - the labeling  $\ell$  is extended by setting  $\ell(v_B) = \text{s1}$ .

**Lemma 2.** Consider a single-epoch execution  $\mathcal{E}$  of `Ouroboros-Chronosmc` with static registration and  $\Delta$ -bounded skew, where  $\Delta$  is the sum of the maximum (local) clock drift and maximum network delay; let  $R$  be the epoch length.

1. Every message sent by an alert party  $P$  in slot  $\text{s1}$  (according to the local time of  $P$ ) will be received by any other alert party  $P'$  by slot  $\text{s1}' \triangleq \text{s1} + \tilde{\Delta}$  for  $\tilde{\Delta} \triangleq 2\Delta$  (according to the local time of  $P'$ ).
2. In particular, we have  $T_{\mathcal{E}} \vdash_{\tilde{\Delta}} w_{\mathcal{E}}$  unless a collision in the responses of the random oracle occurs.

*Proof (sketch).* For the first statement, note that by the assumption `Skew $_{\Delta}$ [s1]`, we know that  $P$  will be executing its (logical) slot  $\text{s1}$  at most  $\Delta$  rounds (by the increase of nominal time) later than  $P'$  executed  $\text{s1}$ ; and if  $P$  sends a message in  $\text{s1}$ , the party  $P$  will receive it at most  $\Delta_{\text{net}}$  activations later after the send event. Note that in the worst case, the send-event happens when the receiving party already has an offset of  $\Delta_{\text{clock}}$  rounds to the sender. Therefore, it will receive it  $\Delta$  rounds later (again, with respect to the nominal time advancement) by the assumption on network delay. Combining these two bounds,  $P'$  will receive the message at latest by slot  $\text{s1}'$  according to its own local clock.

As for the second statement, observe that the properties (i)–(iii) in Definition 7, as well as the requirement that  $\text{range}(\ell) = \{0\} \cup A$ , are satisfied for the same reasons as given in [14, Lemma 6]. The remaining property (iv) is satisfied for  $\tilde{\Delta} \triangleq 2\Delta$  thanks to the first statement of this lemma: Given that an alert party  $P'$  is aware of any block produced by  $P$  for slot  $\text{s1}$ , it will act based upon it and if it creates any block for slot  $\text{s1}'$ , its depth will be strictly larger than the depth of any block created by  $P$  for the slot  $\text{s1}$  by the description of the protocol.  $\square$

To maintain readability, in the following treatment we will omit the (negligible) failure probability caused by random-oracle collisions that are mentioned in the second statement of Lemma 2.

### F.3 Protocol-Induced Distribution of the Characteristic String

Badertscher *et al.* [4] identified the following property of a characteristic string distribution to be of particular interest.

**Definition 9 (The characteristic conditions [4]).** Consider a family of random variables  $W_1, \dots, W_n$  taking values in  $\{0, 1, \perp\}$ . We say that they satisfy the  $(f; \gamma)$ -characteristic conditions if, for each  $k \geq 1$ ,

$$\begin{aligned} \Pr[W_k = \perp \mid W_1, \dots, W_{k-1}] &\geq (1 - f), \\ \Pr[W_k = 0 \mid W_1, \dots, W_{k-1}, W_k \neq \perp] &\geq \gamma, \text{ and hence} \\ \Pr[W_k = 1 \mid W_1, \dots, W_{k-1}, W_k \neq \perp] &\leq 1 - \gamma. \end{aligned}$$

In the expressions above, conditioning on a collection of random variables indicates that the statement is true for any conditioning on the values taken by variables.

The distribution of the characteristic string induced by `Ouroboros-Chronosmc` satisfies  $(f; \gamma)$ -characteristic conditions for parameters recorded in the next lemma.

**Lemma 3 (Protocol-Induced Distribution).** Consider an execution of the protocol  $\text{Ouroboros-Chronos}_{\text{mc}}$  in the single-epoch setting, with static registration and  $\Delta$ -bounded skew. Let  $R$  denote the epoch length and  $f$  be the active-slot coefficient. Let  $\alpha$  (resp.,  $\beta$ ) be a lower bound on the alert stake ratio (resp., participating stake ratio) over the execution. This execution then induces a characteristic string  $W_1, \dots, W_R$  (with each  $W_t \in \{0, 1, \perp\}$ ) satisfying the  $(f; (1-f)^2\alpha)$ -characteristic conditions, and moreover  $\Pr[W_t = \perp \mid W_1, \dots, W_{t-1}] \leq 1 - f\beta$ .

*Proof (sketch).* The lemma can be established by following the same reasoning as in [4, Corollary 2] with respect to logical slots rather than nominal-time rounds.  $\square$

#### F.4 Single-Epoch Security Properties

Any characteristic string that satisfies particular  $(f; \gamma)$ -characteristic conditions does not allow for too large common prefix violations, as proven in [4, Theorem 6] and formally captured by the notion of divergence. We record this result below as Theorem 4, after presenting the necessary formalism in Definitions 10 and 11.

**Definition 10 (Tines, length, and viability [4]).** A path in a fork  $F$  originating at the root is called a tine. For a tine  $t$  we let  $\text{length}(t)$  denote its length, equal to the number of edges on the path. For a vertex  $v$ , we call the length of the tine terminating at  $v$  the depth of  $v$ . For convenience, we overload the notation  $\ell(\cdot)$  so that it applies to tines by defining  $\ell(t) \triangleq \ell(v)$ , where  $v$  is the terminal vertex on the tine  $t$ . We say that a tine  $t$  is  $\Delta$ -viable if  $\text{length}(t) \geq \max_{h+\Delta \leq \ell(t)} \mathbf{d}(h)$ , this maximum extended over all uniquely alert indices  $h$  (appearing  $\Delta$  or more slots before  $\ell(t)$ ). Note that any tine terminating in a uniquely alert vertex is necessarily viable by the  $\Delta$ -monotonicity property.

**Definition 11 (Divergence [26,14]).** Let  $F$  be a  $\Delta$ -fork for a string  $w \in \{0, 1, \perp\}^*$ . For two  $\Delta$ -viable tines  $t$  and  $t'$  of  $F$ , we define the notation  $t/t'$  by the rule

$$t/t' = \text{length}(t) - \text{length}(t \cap t'),$$

where  $t \cap t'$  denotes the common prefix of  $t$  and  $t'$ . Then define the divergence of two viable tines  $t_1$  and  $t_2$  to be the quantity

$$\text{div}(t_1, t_2) = \begin{cases} t_1/t_2 & \text{if } \ell(t_1) < \ell(t_2), \\ t_2/t_1 & \text{if } \ell(t_2) < \ell(t_1), \\ \max(t_1/t_2, t_2/t_1) & \text{if } \ell(t_1) = \ell(t_2). \end{cases}$$

We extend this notation to the fork  $F$  by maximizing over viable tines:  $\text{div}_\Delta(F) \triangleq \max_{t_1, t_2} \text{div}(t_1, t_2)$ , taken over all pairs of  $\Delta$ -viable tines of  $F$ . Finally, we define the  $\Delta$ -divergence of a characteristic string  $w$  to be the maximum over all  $\Delta$ -forks:  $\text{div}_\Delta(w) \triangleq \max_{F \vdash_\Delta w} \text{div}_\Delta(F)$ .

**Theorem 4 ([4, Theorem 6]).** Let  $W = W_1, \dots, W_R$  be a family of random variables, taking values in  $\{0, 1, \perp\}$  and satisfying the  $(f, \gamma)$ -characteristic conditions. If  $\Delta > 0$  and  $\epsilon > 0$  satisfy  $\gamma(1-f)^{\Delta-1} \geq (1+\epsilon)/2$  then

$$\Pr[\text{div}_\Delta(W) \geq k + \Delta] \leq \frac{19R}{\epsilon^4} \exp(-\epsilon^4 k/18).$$

This general statement allows us to translate the properties of the characteristic string distribution induced by an execution of  $\text{Ouroboros-Chronos}_{\text{mc}}$ , as recorded in Lemma 3, into a common-prefix guarantee for the protocol, given below.

**Corollary 2 (Common prefix).** Let  $W = W_1, \dots, W_R$  denote the characteristic string induced by the  $\text{Ouroboros-Chronos}_{\text{mc}}$  protocol in the single-epoch setting with static registration and  $\Delta$ -bounded skew. Let  $R$  be the epoch length and  $f$  the active-slot coefficient. Assume that  $\epsilon > 0$  satisfies

$$\alpha(1-f)^{\tilde{\Delta}+1} \geq (1+\epsilon)/2,$$

where  $\alpha$  is a lower-bound on the alert stake ratio over the execution and  $\tilde{\Delta} = 2\Delta$  is twice the network delay. Then

$$\Pr[\text{div}_{\tilde{\Delta}}(W) \geq k + \tilde{\Delta}] \leq \frac{19R}{\epsilon^4} \exp(-\epsilon^4 k/18),$$

and hence a  $k$ -common-prefix violation occurs with probability at most

$$\bar{\epsilon}_{\text{CP}}(k; R, \Delta, \epsilon) \triangleq \frac{19R}{\epsilon^4} \exp(\tilde{\Delta} - \epsilon^4 k/18).$$

*Proof.* Follows directly from Theorem 4 and Lemma 3, using the first statement of Lemma 2 as a bound on the observed message delivery delay.  $\square$

The remaining Corollaries 3 (resp., 4, 5) below can be established by following the same reasoning as used in the proof of [4, Corollary 4] (resp., [4, Corollary 5], [4, Lemma 11]) with respect to logical slots instead of the nominal time, and using the first statement of Lemma 2 to bound the observed message delivery delay.  $\square$

**Corollary 3 (Chain Growth).** *Let  $W = W_1, \dots, W_R$  denote the characteristic string induced by the Ouroboros-Chronos<sub>mc</sub> protocol in the single-epoch setting with static registration and  $\Delta$ -bounded skew. Let  $R$  be the epoch length and  $f$  the active-slot coefficient. Let  $\alpha, \beta \in [0, 1]$  denote lower bounds on the alert stake ratio and the participating stake ratio over the execution as per Definition 2, and assume that for some  $\epsilon \in (0, 1)$  the parameter  $\alpha$  satisfies*

$$\alpha(1 - f)^{\tilde{\Delta}+1} \geq (1 + \epsilon)/2.$$

where  $\tilde{\Delta} = 2\Delta$  is twice the network delay. Then for

$$s = 48\tilde{\Delta}/(\epsilon\beta f) \quad \text{and} \quad \tau = \beta f/16 \tag{8}$$

we have

$$\Pr[W \text{ admits a } (s, \tau)\text{-CG violation}] \leq \bar{\epsilon}_{\text{CG}}(\tau, s; R, \epsilon) \triangleq \frac{1}{2} s R^2 \exp(-(\epsilon\beta f)^2 s/256).$$

**Corollary 4 (Chain Quality).** *Let  $W = W_1, \dots, W_R$  denote the characteristic string induced by the Ouroboros-Chronos<sub>mc</sub> protocol in the single-epoch setting with static registration and  $\Delta$ -bounded skew. Let  $R$  be the epoch length and  $f$  the active-slot coefficient. Let  $\alpha, \beta \in [0, 1]$  denote lower bounds on the alert stake ratio and the participating stake ratio as per Definition 2, and assume that for some  $\epsilon \in (0, 1)$  the parameter  $\alpha$  satisfies*

$$\alpha(1 - f)^{\tilde{\Delta}+1} \geq (1 + \epsilon)/2.$$

where  $\tilde{\Delta} = 2\Delta$  is twice the network delay. Then for

$$k = 48\tilde{\Delta}/(\epsilon\beta f) \quad \text{and} \quad \mu = \epsilon\beta f/16$$

we have

$$\Pr[W \text{ admits a } (\mu, k)\text{-CQ violation}] \leq \bar{\epsilon}_{\text{CQ}}(\mu, k; R, \epsilon) \triangleq \frac{1}{2} k R^2 \exp(-(\epsilon\beta f)^2 k/256).$$

**Corollary 5 (Existential Chain Quality).** *Let  $W = W_1, \dots, W_r$  denote the characteristic string induced by the protocol Ouroboros-Chronos in the single-epoch setting with static registration and  $\Delta$ -bounded skew. Let  $R$  be the epoch length and  $f$  the active-slot coefficient. Let  $\alpha, \beta \in [0, 1]$  denote lower bounds on the alert stake ratio and the participating stake ratio over the execution as per Definition 2, and assume that for some  $\epsilon \in (0, 1)$  the parameter  $\alpha$  satisfies*

$$\alpha(1 - f)^{\tilde{\Delta}+1} \geq (1 + \epsilon)/2,$$

where  $\tilde{\Delta} = 2\Delta$  is twice the network delay. Then for  $s \geq 12\tilde{\Delta}/(\epsilon\beta f)$ ,

$$\Pr[W \text{ admits a } s\text{-}\exists\text{CQ violation}] \leq \bar{\epsilon}_{\exists\text{CQ}}(s; r, \epsilon) = r^2(s + 1) \exp(-(\epsilon\beta f)^2 s/64).$$

## F.5 Switching to maxvalid-bg

To capture the security of the full protocol Ouroboros Chronos with the chain-selection rule maxvalid-bg given in Section C.8, the bounds above need to be adjusted by an additional term that reflects the probability that the new chain selection rule deviates from the easier-to-analyze rule maxvalid-mc. This error term was quantified by Theorem 2 in [4] and this quantification translates directly into our setting.

**Corollary 6 (Security of maxvalid-bg).** *Consider the protocol Ouroboros-Chronos (with maxvalid-bg), executed in the same setting and under the same assumptions as in Corollaries 2–5 above. If the maxvalid-bg parameters,  $k$  and  $s$ , satisfy*

$$k > 192\tilde{\Delta}/(\epsilon\beta) \quad \text{and} \quad R/6 \geq s = k/(4f) \geq 48\tilde{\Delta}/(\epsilon\beta f)$$

*then the guarantees given in Corollaries 2–5 for common prefix, chain growth, chain quality and existential chain quality are also valid for Ouroboros-Chronos except for an additional error probability*

$$\bar{\epsilon}_{mv} \triangleq \exp(\ln R - \Omega(k)) + \bar{\epsilon}_{CG}(\beta f/16, k/(4f)) + \bar{\epsilon}_{\exists CQ}(k/(4f)) + \bar{\epsilon}_{CP}(k\beta/64), \quad (9)$$

where the subscript “mv” stands for “maxvalid”.

*Proof (sketch).* The corollary follows by the same reasoning as Theorem 2 in [4], again by using the first statement of Lemma 2 to bound the observed message delivery delay.  $\square$

**Putting things together.** The proof of Theorem 1 now follows by combining Corollaries 2, 3, 4, 5 and 6.

## G Analysis of the Synchronization Procedure

### G.1 The Proof of Fact 1

*Proof (of Fact 1).* Without loss of generality order the pairs so that  $a_i \leq a_j$  for  $i < j$  (and note that this does not necessarily imply  $b_i \leq b_j$  for  $i < j$ ). Now we have

$$\begin{aligned} a_{\lceil n/2 \rceil} - \Delta &\stackrel{(a)}{\leq} \min \{b_i : \lceil n/2 \rceil \leq i \leq n\} \stackrel{(b)}{\leq} \text{med}((b_i)_{i=1}^n) \\ &\stackrel{(c)}{\leq} \max \{b_i : 1 \leq i \leq \lceil n/2 \rceil\} \stackrel{(d)}{\leq} a_{\lceil n/2 \rceil} + \Delta, \end{aligned}$$

where inequalities (a) and (d) follow from the assumption of the lemma and the assumed ordering of the values  $a_i$ , inequalities (b) and (c) follow from the definition of med. The proof is concluded by observing that  $\text{med}((a_i)_{i=1}^n) = a_{\lceil n/2 \rceil}$  by definition.  $\square$

### G.2 SyncProc maintains Skew $_{\Delta}$

**Lemma 4 (SyncProc maintains Skew $_{\Delta}$ ).** *Consider an execution of the full protocol Ouroboros-Chronos over a lifetime of  $L = ER$  slots, where  $R$  is the epoch length. Let  $\Delta_{\text{net}}$  be the upper bound on message delay enforced by  $\mathcal{F}_{N\text{-}MC}$  and let  $\Delta_{\text{clock}}$  be the maximal drift enforced by  $\mathcal{G}_{\text{IMPERF}LCLOCK}$  and let  $\Delta = \Delta_{\text{clock}} + \Delta_{\text{net}}$ ; and let  $\mathbf{s1} \geq 1$  be the last slot of some epoch  $\mathbf{ep} \geq 1$ , i.e., such that  $\mathbf{s1} \bmod R = 0$ . If the properties  $CG(\tau_{CG}, R/3)$  and  $CP(\tau_{CG}R/3)$  for  $\tau_{CG}$  as in Theorem 1 are not violated during the execution up to slot  $\mathbf{s1}$ , then the predicate  $\text{Skew}_{\Delta}[\mathbf{s1}']$  is satisfied for  $\mathbf{s1}' = \mathbf{s1} + 1$  onward to the next synchronization slot.*

*Proof.* We first establish two intermediate claims under the lemma assumptions:

- (i) All alert parties use the same set of synchronization beacons in their execution of the procedure SyncProc between epochs  $\mathbf{ep}$  and  $\mathbf{ep} + 1$ , formally  $\mathcal{S}_j^{\mathbf{P}_1} = \mathcal{S}_j^{\mathbf{P}_2}$  for any two parties  $\mathbf{P}_1, \mathbf{P}_2$  that are alert in the  $j$ -th synchronization slot.
- (ii) For any fixed beacon  $\mathbf{SB} \in \mathcal{S}_j^{\mathbf{P}_1} = \mathcal{S}_j^{\mathbf{P}_2}$ , the quantity

$$\mu(\mathbf{P}_i, \mathbf{SB}) \triangleq \text{Skew}^{\mathbf{P}_i}[\mathbf{s1}] + \text{slotnum}(\mathbf{SB}) - \mathbf{P}_i.\text{Timestamp}(\mathbf{SB})$$

will differ by at most  $\Delta$  between any two alert parties  $\mathbf{P}_1$  and  $\mathbf{P}_2$ .

To see (i), note that the set  $\mathcal{S}_j^P$  is constructed by the party  $P$  by collecting all beacons  $\mathbf{SB}$  that report a slot number  $\text{slotnum}(\mathbf{SB})$  within  $[(i-1) \cdot R + 1, \dots, (i-1) \cdot R + R/6]$  (the preceding synchronization interval) and which are included in a block of the adopted chain  $P.\mathcal{C}_{\text{loc}}$  of  $P$  up to slot  $(i-1) \cdot R + 2R/3$ . Based on the chain growth property, the chain  $\mathcal{C}_{\text{loc}}$  contain as least  $\tau_{\text{CG}}R/3$  blocks in the last  $R/3$  slots, and by the common prefix property, the chains are identical up to slot  $(i-1) \cdot R + 2R/3$ .

Now observe that

$$\mu(P_i, \mathbf{SB}) = (\text{slotnum}(\mathbf{SB}) - t) - \theta_{P_i, \mathbf{SB}} - \delta_{P_i, \mathbf{SB}}, \quad (10)$$

where  $t$  is the nominal time at the point of the execution where  $\mathbf{SB}$  was sent, and  $\delta_{P_i, \mathbf{SB}} \in [\Delta_{\text{net}}]$  is the number of rounds that  $\mathbf{SB}$  was delayed in its transit from  $\tilde{P}$  to  $P_i$ , i.e., where one round is counted per local clock update of  $P_i$ , and where  $\theta_{P_i, \mathbf{SB}} \in [\Delta_{\text{clock}}]$  is the drift of  $P_i$  from the nominal time  $t$  at the point of the execution when  $\mathbf{SB}$  was sent. Clearly, equation (10) establishes (ii). Note that (10) holds even if the sender  $\tilde{P}$  of  $\mathbf{SB}$  was corrupted when sending it (owing to the network model); for alert parties  $\tilde{P}$  the first bracket in (10) is equal to  $\text{Skew}^{\tilde{P}}[\text{slotnum}(\mathbf{SB})]$  which includes any drift from the sender to the nominal time by definition.

Given the above properties, we invoke Fact 1 for the tuples

$$(\mu(P_1, \mathbf{SB}))_{\mathbf{SB} \in \mathcal{S}_j^{P_1}} \text{ and } (\mu(P_2, \mathbf{SB}))_{\mathbf{SB} \in \mathcal{S}_j^{P_2}}$$

for two arbitrary alert parties  $P_1$  and  $P_2$ . By property (i) both tuples are of the same size, and by property (ii) one can see that the update mechanism of Chronos will yield close time stamps because the preconditions of Fact 1 are fulfilled by some  $\Delta' \leq \Delta$ . Therefore we obtain

$$\left| \text{med} \left( (\mu(P_1, \mathbf{SB}))_{\mathbf{SB} \in \mathcal{S}_j^{P_1}} \right) - \text{med} \left( (\mu(P_2, \mathbf{SB}))_{\mathbf{SB} \in \mathcal{S}_j^{P_2}} \right) \right| \leq \Delta'$$

and further, for  $i \in \{1, 2\}$ ,  $\text{med} \left( (\mu(P_i, \mathbf{SB}))_{\mathbf{SB} \in \mathcal{S}_j^{P_i}} \right) = \text{Skew}^{P_i}[\mathfrak{s}1 + 1]$  by the update mechanism of Chronos. Lemma 4 finally follows by observing that the maximal drift of a party in the execution of the system is an absolute term, and thus the quantity  $\Delta_{\text{clock}} - \max\{\theta_{P_i, \mathbf{SB}} \mid \mathbf{SB} \in \mathcal{S}_j^{P_i}\}$  is the remaining additional term (measured in number of rounds) that the party  $P_i$  can be pushed forward (and hence increasing the actual skew  $\Delta'$ ) in the execution without the baseline (and therefore everyone) moving forward. We see from the computation in equation (10) that the the skew will never exceed  $\Delta$  and can thus establish  $\text{Skew}_{\Delta}[\mathfrak{s}1']$  for  $\mathfrak{s}1' = \mathfrak{s}1 + 1$  and onward up to the next synchronization slot  $j + 1$ .  $\square$

### G.3 Bounded shift

**Lemma 5 (Bounded shift).** *Consider an execution of the full protocol **Ouroboros-Chronos** over a lifetime of  $L = ER$  slots, where  $R$  is the epoch length. Let  $\Delta_{\text{net}}$  be the upper bound on message delay enforced by  $\mathcal{F}_{\text{N-MC}}$  and let  $\Delta_{\text{clock}}$  be the maximal drift enforced by  $\mathcal{G}_{\text{IMPERFLOCK}}$  and let  $\Delta = \Delta_{\text{clock}} + \Delta_{\text{net}}$ . Further assume  $\tilde{\Delta} \triangleq 2\Delta \leq R/6$ . Let  $\mathfrak{s}1 \geq 1$  be the last slot of some epoch  $\mathfrak{ep} \geq 1$ , i.e., such that  $\mathfrak{s}1 \bmod R = 0$ , and assume that  $\text{Skew}_{\Delta}[\mathfrak{s}1']$  is satisfied for all slots in epoch  $\mathfrak{ep}$ . Let  $\alpha \in [0, 1]$  denote a lower bound on the alert ratio and participating ratio throughout the execution. If for some  $\epsilon \in (0, 1)$  we have  $\alpha \cdot (1 - f) \geq (1 + \epsilon)/2$ , and if the property  $\exists\text{CQ}(R/3)$  is not violated during the execution up to slot  $\mathfrak{s}1$ , then in any invocation of **SyncProc** by an alert party during  $\mathfrak{s}1$ , the local variable **shift** computed on line 17 will satisfy  $|\text{shift}| \leq 2\Delta$ , except with error probability  $\exp(\ln L - \Omega(R))$  over the whole execution.*

*Proof.* We first show that the set  $\mathcal{S}_i^P$  used by any alert party  $P$  contains all beacon messages produced by alert parties in the last synchronization interval. This follows by observing that all synchronization beacons produced by alert parties in (their) slots  $[(i-1) \cdot R + 1, \dots, (i-1) \cdot R + R/6]$  will be delivered to other alert parties by the slot  $(i-1) \cdot R + R/3$  by the assumption  $\tilde{\Delta} \leq R/6$  and by the first statement of Lemma 2. Moreover, by the  $\exists\text{CQ}$  property, the chain  $\mathcal{C}_{\text{loc}}$  of any alert party  $P$  will contain at least one block created by an alert party over the slot interval  $[(i-1) \cdot R + R/3 + 1, \dots, (i-1) \cdot R + 2R/3]$ . When such an alertly-created block is included, it will contain all the synchronization beacons produced by alert parties for slots  $[(i-1) \cdot R + 1, \dots, (i-1) \cdot R + R/6]$  that were not included yet.

Next, in light of the lower bound on alert stake ratio, a majority of synchronizing beacons in  $\mathcal{S}_i^P$  will be alertly-generated, except with error probability  $\exp(\ln L - \Omega(R))$ . Therefore, if this error does not occur, due

to the use of median in the definition of `shift` there exist alert parties  $P_1, P_2$ , which produced synchronization beacons  $SB_1, SB_2 \in \mathcal{S}_i^P$  respectively, such that

$$\text{Skew}^{P_1}[\text{slotnum}(SB_1)] - \Delta \stackrel{(a)}{\leq} \text{Skew}^P[\mathbf{s1}] + \text{shift} \stackrel{(b)}{\leq} \text{Skew}^{P_2}[\text{slotnum}(SB_2)] .$$

The inequalities (a) and (b) follow again from the core computation in equation (10), this time considering  $\mu(P, SB_1)$  and  $\mu(P, SB_2)$  and the bounds  $\delta_{P, SB_1} + \theta_{P_i, SB_1} \leq \Delta$  and  $\delta_{P, SB_2}, \theta_{P_i, SB_2} \geq 0$ , respectively. The proof is now concluded by observing that for both  $i \in \{1, 2\}$ , we have

$$\left| \text{Skew}^{P_i}[\text{slotnum}(SB_i)] - \text{Skew}^P[\mathbf{s1}] \right| \leq \Delta$$

thanks to the assumption  $\text{Skew}_\Delta[\mathbf{s1}]$ . □

## H Security for the Full Execution with Static Registration

The following theorem is a formal statement corresponding to the informal Theorem 2.

**Theorem 5 (Full-execution security with static registration).** *Consider the execution of Ouroboros-Chronos with adversary  $\mathcal{A}$  and environment  $\mathcal{Z}$  in the setting with static registration. Let  $f$  be the active-slot coefficient, let  $\Delta$  be the upper bound on the sum of the maximum network delay and maximum local clock drifts, and let  $\tilde{\Delta} \triangleq 2\Delta$ . Let  $\alpha, \beta \in [0, 1]$  denote a lower bound on the alert and participating stake ratios throughout the whole execution, respectively. Let  $R$  and  $L$  denote the epoch length and the total lifetime of the system (in slots), and let  $Q$  be the total number of queries issued to  $\mathcal{G}_{\text{RO}}$ . If the assumptions (5) and (6) are satisfied, then Ouroboros-Chronos achieves the same guarantees for common prefix (resp. chain growth, chain quality, existential chain quality) as given in Theorem 1 (with  $L$  replacing  $R$  as execution length) except with an additional error probability of*

$$QL \cdot (\bar{\epsilon}_{\text{CG}}(\tau_{\text{CG}}, R/3; R, \epsilon) + \bar{\epsilon}_{\text{CP}}(\tau_{\text{CG}}R/3; R, \Delta, \epsilon) + \bar{\epsilon}_{\exists\text{CQ}}(R/3; R, \epsilon)) , \quad (11)$$

where  $\tau_{\text{CG}} = \beta f/16$ . If  $R \geq 144\tilde{\Delta}/\epsilon\beta f$  then this term can be upper-bounded by

$$\epsilon_{\text{lift}} \triangleq QL \cdot \left[ R^3 \cdot \exp\left(-\frac{(\epsilon\beta f)^2 R}{768}\right) + \frac{19R}{\epsilon^4} \cdot \exp\left(\tilde{\Delta} - \frac{\epsilon^4 \tau_{\text{CG}} R}{54}\right) + 3\bar{\epsilon}_{\text{mv}} \right] . \quad (12)$$

For all  $\mathbf{p} \in \{\text{CP}, \text{CG}, \exists\text{CQ}, \text{CQ}\}$ , we denote the obtained counterparts of the single-epoch error terms  $\bar{\epsilon}_{\mathbf{p}}$  for the full execution with static registration by  $\epsilon_{\mathbf{p}}$ .

*Proof (of Theorem 5, sketch).* When moving from the single-epoch setting to a setting with several epochs, the following aspects need to be considered:

- **Stake distribution updates.** The stake distribution used for sampling slot leaders changes in every epoch. In Ouroboros-Chronos the distribution used for sampling in epoch  $\mathbf{ep}$  is set to be the stake distribution recorded on the blockchain up to the last block of the epoch  $\mathbf{ep} - 2$ .
- **Randomness updates.** Every epoch needs new public randomness to be used for the private leader election process based on the above distribution. For epoch  $\mathbf{ep}$ , this randomness is obtained by hashing together VRF-outputs put into blocks in epoch  $\mathbf{ep} - 1$  by their creators. More precisely, the protocol hashes together these values from the blocks in the first  $2R/3$  slots of epoch  $\mathbf{ep} - 1$  (out of its  $R$  slots).
- **Resynchronization.** All alert parties execute the resynchronization procedure `SyncProc` in the last slot of every epoch.

This proof partly follows the treatment in Section 5 of [14] to argue about stake distribution and randomness updates, and hence we only sketch the reasoning for these parts, adding a discussion of the resynchronization. The proof has an inductive structure over the epochs of the execution.

First, note that in the first epoch, we have both perfect epoch randomness, and the property  $\text{Skew}_\Delta[\mathbf{s1}]$  is satisfied for all slots  $\mathbf{s1}$  in this epoch, by definition of the functionality  $\mathcal{F}_{\text{INIT}}$  and  $\mathcal{G}_{\text{IMPERFLCLOCK}}$ . More precisely, they ensure that  $\text{Skew}_\Delta[\mathbf{s1}]$  for all slots in the first epoch (where  $\Delta$  is the bound on the sum of all incurred delays and clock inaccuracies) and all alert parties advance their local clock at least at the pace of



the baseline. Given that, we can apply Theorem 1 to the first epoch and obtain the bounds for CP, CG and CQ it provides.

For the induction step, Lemma 4 shows that the properties  $\text{CG}(\tau_{\text{CG}}, R/3)$  and  $\text{CP}(\tau_{\text{CG}}R/3)$  satisfied in epoch  $\text{ep}$  imply  $\text{Skew}_{\Delta}[\text{s1}]$  for all slots in epoch  $\text{ep} + 1$ . Furthermore,  $\text{CG}(\tau_{\text{CG}}, R/3)$  and  $\text{CP}(\tau_{\text{CG}}R/3)$  during the first  $R/3$  slots of epoch  $\text{ep}$  also imply that each alert player’s chain grows by at least  $\tau R/3$  blocks and after these slots, all alert players agree on the stake distribution at the end of epoch  $\text{ep} - 1$ . Moreover,  $\exists \text{CQ}(R/3)$  implies that during the second  $R/3$  slots of epoch  $\text{ep}$ , each alert player’s chain contains at least one honest block. Hence the randomness that will be derived for epoch  $\text{ep} + 1$  will be influenced by at least one honest VRF-output chosen *after* the stake distribution for  $\text{ep} + 1$  is fixed. Finally,  $\text{CG}(\tau_{\text{CG}}, R/3)$  and  $\text{CP}(\tau_{\text{CG}}R/3)$  during the last  $R/3$  slots of epoch  $\text{ep}$  imply that each alert player’s chain grows by at least  $\tau R/3$  blocks and therefore after these slots, all alert players agree on the randomness for the epoch  $\text{ep} + 1$ .

Exactly as in [14,4], we need to additionally account for the limited amount of grinding that the adversary can achieve by deciding whether to include blocks (with VRF outputs) in slots where he is a slot leader. This can be crudely upper-bounded by limiting the number of queries to the random oracle that the adversary makes, adding an additional quantity  $Q$  into our bound.

Now, having established  $\text{Skew}_{\Delta}[\text{s1}]$  for the slots in epoch  $\text{ep} + 1$  as well as accounted for the quality of the randomness used in epoch  $\text{ep} + 1$ , we can again invoke Theorem 1 to include  $\text{ep} + 1$  to obtain guarantees on CP, CG, and CQ and complete the induction step.

Finally, the bound (12) is obtained by instantiating (11) with the concrete bounds of Theorem 1.  $\square$

## I Analysis of Joining

**Lemma 6.** *Consider an execution of the full protocol `Ouroboros-Chronos` and let  $P_{\text{join}}$  be a party joining the protocol execution at time  $t_{\text{join}} > 0$  that retains its access to all resources during its joining procedure `JoinProc` (cf. Fig. 2). Let  $t \in (t_{\text{join}} + t_{\text{off}}, t_{\text{join}} + t_{\text{off}} + t_{\text{minSync}} + t_{\text{stable}} + 1]$  be the (nominal) time at some point in the execution in which  $P_{\text{join}}$  is in Phase C or D of its joining procedure and let  $C_{\text{join}}$  denote the chain held by  $P_{\text{join}}$  at that point. Let  $C_{\text{alert}}$  denote a chain held by any alert party  $P_{\text{alert}}$  at some point in the execution where time  $t' \triangleq t - \Delta$ . Then under the assumptions of Theorem 5 and assuming no violations of  $\text{CP}(k\beta/64)$ ,  $\exists \text{CQ}(s)$ , and  $\text{CG}(\tau_{\text{CG}}, s)$  until the end of the joining procedure (for the parameters  $k$  and  $s$  of `maxvalid-bg`), we have  $C_{\text{alert}}^{\lceil k} \preceq C_{\text{join}}$  except with error probability  $\exp(\ln L - \Omega(R))$  over the whole execution.*

*Proof (sketch).* Notice that the chain-selection procedure `maxvalid-bg` given in Fig. C.8 does not involve the local time `localTime` of the party executing it. Therefore,  $P_{\text{join}}$  and  $P_{\text{alert}}$  would do the same chain-selection decisions in their `maxvalid-bg` procedures, if presented with the same inputs. The only (but crucial) difference in their chain-adoption behavior comes from the fact that  $P_{\text{alert}}$  has local clock that satisfies the  $\text{Skew}_{\Delta}$  predicate, and based on this local clock the party removes from consideration all received chains that contain blocks from its logical future (with timestamp larger than its local time), this is done on line 4 of procedure `IsValidChain` in Fig. C.5. Of course,  $P_{\text{join}}$  does no such filtering as it does not possess reliable local time information yet.

To see the implications of this difference, we consider the concept of a *virtual execution* for  $P_{\text{join}}$  introduced in [4]. This is an artificial random experiment that consists of the execution of the protocol with an additional (“*virtual*”) party  $P_{\text{virt}}$  that participates from the beginning, is always alert, but commands no stake and hence is passive. Moreover, starting from the point of execution where  $P_{\text{join}}$  joins the system, the virtual party advances exactly like  $P_{\text{join}}$  and receives the same messages in the same slots and order as  $P_{\text{join}}$ .

The only case when  $P_{\text{join}}$  may adopt as its local chain a chain  $C_{\text{join}}$  that  $P_{\text{virt}}$  does not adopt over the chain it is currently holding (call it  $C_{\text{virt}}$ ) is if  $C_{\text{join}}$  contains an adversarially-created suffix of future blocks such that it dominates  $C_{\text{virt}}$  based on Condition A in `maxvalid-bg`. (As proved in [4, Theorem 2], the adversary is not capable of creating a chain that would dominate an alert chain according to Condition B under the assumptions of the lemma, except for a global bad event with probability  $\exp(\ln L - \Omega(R))$ .) However, Condition A is only applied if  $C_{\text{virt}}^{\lceil k} \preceq C_{\text{join}}$  so this must have been the case when  $P_{\text{join}}$  adopted  $C_{\text{join}}$  prior to the point of the execution at time  $t$  we are considering in the lemma statement. Moreover, since  $P_{\text{join}}$  is still holding  $C_{\text{join}}$  at the point under consideration at time  $t$ , it means that since it adopted it,  $P_{\text{virt}}$  has not received

any chain  $C'_{\text{virt}}$  that would violate  $C'_{\text{virt}} \preceq C_{\text{join}}$ , as in that case also  $P_{\text{join}}$  would receive and adopt it (as follows by inspection of `maxvalid-bg`). Therefore, the chain  $C'_{\text{virt}}$  that  $P_{\text{virt}}$  holds at the point under consideration at time  $t$  satisfies  $C'_{\text{virt}} \preceq C_{\text{join}}$ . Finally, if any alert party  $P_{\text{alert}}$  held at nominal time  $t'$  a chain  $C_{\text{alert}}$  that would violate  $C_{\text{alert}} \preceq C_{\text{join}}$ , by our assumptions on the network delay and clock drift, it would be delivered and considered by  $P_{\text{virt}}$  at any point in the execution where nominal time is  $t$  and hence adopted, concluding the proof.  $\square$

**Lemma 7.** *Consider an execution of the full protocol `Ouroboros-Chronos` and let  $P_{\text{join}}$  be a party joining the protocol execution at time  $t_{\text{join}} > 0$  that retains its access to all resources during its joining procedure `JoinProc` (cf. Fig. 2). Under the assumptions of Theorem 5 and Lemma 6, and assuming no violations of  $\text{CG}(\tau_{\text{CG}}, R/3)$ ,  $\text{CP}(\tau_{\text{CG}}R/3)$ , and  $\exists\text{CQ}(R/3)$  until the end of the joining procedure, we have the following except with error probability  $\exp(\ln L - \Omega(R))$  over the whole execution:*

- (a) *The index value  $i^*$  determined on line 32 of its joining procedure `JoinProc` satisfies  $i^* \geq 1$ .*
- (b) *For all values of  $i \geq i^*$  processed in the iteration on lines 34–52 we have  $\mathcal{S}_i^{\text{P}_{\text{join}}} = \mathcal{S}_i^{\text{P}_{\text{alert}}}$ , where  $\mathcal{S}_i^{\text{P}_{\text{join}}}$  is the set of synchronization beacons determined by  $P_{\text{join}}$  on line 35 and  $\mathcal{S}_i^{\text{P}_{\text{alert}}}$  is the set of synchronization beacons determined by any alert party  $P_{\text{alert}}$  on line 6 of its procedure `SyncProc` for the same  $i$ .*
- (c) *For all values of  $i \geq i^*$  processed in the iteration on lines 34–52 and for any fixed beacon  $\text{SB} \in \mathcal{S}_i^{\text{P}_{\text{join}}} = \mathcal{S}_i^{\text{P}_{\text{alert}}}$ , the quantity  $\mu(P, \text{SB})$ , which is defined to be  $\text{Skew}^P[\text{s1}] + \text{slotnum}(\text{SB}) - \text{P.Timestamp}(\text{SB})$ , will differ by at most  $\Delta$  between the two parties  $P \in \{P_{\text{join}}, P_{\text{alert}}\}$ , and furthermore, when the joining party becomes alert  $\text{Skew}_{\Delta}[\text{s1}]$  is satisfied for any slot in which the joining party is considered alert.*

*Proof (sketch).* We first show the claim (a) that condition  $i^* \geq 1$  on line 32 of `JoinProc` will be satisfied for  $P_{\text{join}}$ . Informally speaking, this means that while  $P_{\text{join}}$  executed Phase C of its joining procedure `JoinProc` (recall that line 32 is only executed after that), it has observed at least one full synchronization interval  $I_{\text{sync}}(i^*)$  that started at least  $t_{\text{pre}}$  rounds after the beginning of Phase C; and has recorded timestamps (in its data structure `Timestamp`) for all synchronization beacons  $\text{SB}$  recorded in its local chain and coming from  $I_{\text{sync}}(i)$  according to their included logical slot numbers. Also recall that when we talk about an interval  $d$  of rounds that a party locally executes, then this relates to the number of rounds all other parties have executed by an additive offset  $\delta \leq \Delta_{\text{clock}}$  and we choose the all intervals in this proof in such a way that it encompasses this small offset.

To proceed, let us split Phase C into two consecutive, non-overlapping Phases  $C_{\text{sync}}$  and  $C_{\text{stable}}$  consisting of  $t_{\text{minSync}}$  and  $t_{\text{stable}}$  rounds, respectively. Let  $t_{\text{start}}^{(j)}$  denote the nominal time in which it happens for the first time that an alert party enters the synchronization interval  $I_{\text{sync}}(j)$  according to its local clock (i.e., enters the logical slot  $(j-1)R+1$ ). Then for all  $j \geq 1$  we have

$$t_{\text{start}}^{(j+1)} - t_{\text{start}}^{(j)} \leq R + 3\Delta \leq 13R/12 \quad (13)$$

thanks to the fact that there is a synchronization interval starting at the first slot of every  $R$ -slot epoch, the bound of Lemma 5 and the assumptions  $\text{Skew}_{\Delta}$  and (6).

Let now  $i^*$  denote the minimal  $i$  such that  $t_{\text{start}}^{(i)} \geq t_{\text{join}} + t_{\text{off}} + t_{\text{pre}}$ , i.e.,  $t_{\text{start}}^{(i^*)}$  occurs at least  $t_{\text{pre}}$  rounds after the beginning of  $P_{\text{join}}$ 's Phase C. According to (13) we have

$$t_{\text{start}}^{(i^*)} \leq t_{\text{join}} + t_{\text{off}} + t_{\text{pre}} + 13R/12 \leq t_{\text{join}} + t_{\text{off}} + t_{\text{minSync}} - R/6$$

by the values of  $t_{\text{pre}}$  and  $t_{\text{minSync}}$  (cf. Table 1). Therefore  $t_{\text{start}}^{(i^*)}$  is guaranteed to occur at least  $R/6$  rounds before the end of Phase  $C_{\text{sync}}$  for  $P_{\text{join}}$ .

We now argue that  $i^*$  will satisfy the condition on line 32 of `JoinProc`. This follows as by round  $t_{\text{start}}^{(i^*)}$ ,  $P_{\text{join}}$  has been already recording the timestamps of all received synchronization beacons for  $t_{\text{pre}}$  rounds, and hence, has recorded into its `Timestamp` data structure the timestamps of all beacons that (according to the logical slot number they contain) belong to the synchronization interval  $I_{\text{sync}}[i^*]$  — either by receiving the beacon directly or observing it as included in a blockchain block. This is exactly what is needed for  $i^*$  to pass the test on line 32 of `JoinProc`. Note that the adversary cannot create valid beacons logically belonging to  $I_{\text{sync}}[i]$

before the start of  $P_{\text{join}}$ 's Phase C, as before that point the epoch randomness necessary for creating valid synchronization beacons for this synchronization interval is still completely unpredictable (thanks to the choice of  $t_{\text{pre}}$ ).

Moving to claim (b), we first establish it for  $i^*$ . This can be argued in a similar way as the validity of item (i) in the proof of Lemma 4: the set  $\mathcal{S}_{i^*}^P$  is for both  $P \in \{P_{\text{join}}, P_{\text{alert}}\}$  constructed by collecting all beacons SB (satisfying certain conditions on the reported slot number) from the adopted chain  $P.C_{\text{loc}}$  of  $P$  up to slot  $(i^* - 1) \cdot R + 2R/3$ . As observed above, the synchronization interval  $I_{\text{sync}}(i^*)$  will start (from the perspective of the first alert party) at least  $R/6$  rounds before the end of  $P_{\text{join}}$ 's Phase  $C_{\text{sync}}$ , and hence will also end (for this alert party) before the end of Phase  $C_{\text{sync}}$ . Therefore, it will be followed by  $t_{\text{stable}} = R$  rounds of Phase  $C_{\text{stable}}$ , and the relevant beacons will be collected from up to round  $2R/3$  of Phase  $C_{\text{stable}}$  (to account for the potential skew of other alert parties and delays plus drifts, as in (13)). Assuming  $\text{CG}(\tau_{\text{CG}}, R/3)$  and  $R \geq 3k\tau_{\text{CG}}$  (which follows from (6)), we get that the chain held by  $P_{\text{alert}}$  grows by at least  $k$  blocks during the last  $R/3$  slots, and is hence identical to the chain held by  $P_{\text{join}}$  up to slot  $(i^* - 1) \cdot R + 2R/3$  by Lemma 6, resulting in  $\mathcal{S}_{i^*}^{P_{\text{join}}} = \mathcal{S}_{i^*}^{P_{\text{alert}}}$ . The above reasoning applies identically also to all following values of  $i \geq i^*$  that the iteration on lines 34–52 considers.

Finally, claim (c) follows by similar arguments as in the proof of Lemma 4 (relying on the network and clock model guarantees).  $\square$

## J Full Security in the Dynamic Availability Setting

The following theorem is a formal statement corresponding to the informal Theorem 3.

**Theorem 6 (Dynamic availability).** *Consider an execution of the full protocol Ouroboros-Chronos in the dynamic-availability setting. Under the assumptions of Theorem 5 and Lemma 7, Ouroboros-Chronos achieves the same guarantees for common prefix (resp. chain growth, chain quality, existential chain quality) as given in Theorem 5 except for the negligible additional error probability*

$$\epsilon_{\text{DA}} \triangleq \epsilon_{\text{CP}} \max \left\{ \frac{k\beta}{64}, \frac{\tau_{\text{CG}}R}{3} \right\} + \epsilon_{\text{CG}}(\tau_{\text{CG}}, s) + \epsilon_{\exists\text{CQ}}(\tau_{\text{CG}}R/3) + e^{\ln L - \Omega(R)}.$$

*Proof (of Theorem 6, sketch).* The theorem follows by considering each of the new situations that occur when honest parties lose and regain some of their resources. We sketch these considerations below.

If an alert party briefly loses access to its random oracle, it will keep the synchronized status, and start caching all network messages and advancing its local clock “blindly” by 1 slot per tick of  $\mathcal{G}_{\text{IMPERFLCLOCK}}$ . Hence, it will not violate the  $\text{Skew}_{\Delta}$  invariant until it first reaches a synchronization slot, and is able to become alert again immediately upon regaining access to  $\mathcal{G}_{\text{RO}}$ . However, once it reaches a synchronization slot, it declares itself desynchronized, hence not affecting  $\text{Skew}_{\Delta}$  anymore. Similarly, if an honest party loses access to its clock or its network, it immediately becomes desynchronized.

Desynchronized parties maintain this status until they regain all resources, at which point they run the joining procedure  $\text{JoinProc}$  analyzed in Section 6.5, just like newly joining parties. The claims (b) and (c) of Lemma 7 guarantee that upon completion of this procedure, when the party declares itself synchronized again, it will indeed satisfy the invariant  $\text{Skew}_{\Delta}$  except with a negligible error probability: this follows again by invoking Fact 1 in the same way as done for synchronized parties in Lemma 4 based on the claims (i) and (ii) from its proof.

Finally, note that the proof of Theorem 1 relies on the martingale analysis from Appendix F, which was designed in [4] exactly for the purpose of carrying over to the dynamic availability setting that allows the environment to adjust the stake ratios of alert and active parties adaptively during the execution of the protocol.  $\square$

## K Connection between Logical Time and Nominal Time

We first state and prove the general lemma:

**Lemma 8 (Nominal vs. logical time growth).** *Consider an execution of the full protocol Ouroboros-Chronos in the dynamic-availability setting, let  $P$  be a party that is synchronized between (and including) slots  $s_1$  and  $s_1'$ , let  $t$  and  $t'$  be the nominal times when  $P$  enters slot  $s_1$  and  $s_1'$  for the first time, respectively. Denote by  $\delta s_1$  and  $\delta t$  the respective differences  $|s_1' - s_1|$  and  $|t' - t|$ . Define the quantity  $\tau_{\text{TG}} \triangleq 1 - (96\Delta + R\epsilon\beta f)/(48R)$ . Then, under the assumptions of Theorem 6, we have  $\delta s_1 \geq \tau_{\text{TG}} \cdot \delta t$  whenever  $\delta t \geq 48\tilde{\Delta}/(\epsilon\beta f)$  (where  $\tilde{\Delta} = 2\Delta$ ).*

*Proof (of Lemma 8).* The lemma follows directly from Lemma 5 carried over to the dynamic-availability setting. In particular, the skew that the adversary can apply in every sequence when nominal time increases by  $R$  is at most  $2\Delta$  since no more synchronization slots can occur where the synchronized parties adjust their local time-stamps. In between, all alert parties proceed at least at the baseline speed, which is used to define the nominal time. Given that the interval under consideration could start right at a synchronization slot of alert party  $P$ , we need to incorporate an additional offset of  $2\Delta$  giving a total shift of at most  $2\Delta \cdot \delta t/R + 2\Delta$ . Relative to  $\delta t$ , this shift can be expressed as  $(2\Delta/R + x)\delta t$  for some  $x > 0$  as long as  $\delta t \geq 2\Delta/x$ . For the sake of concreteness, we pick  $x = (\epsilon\beta f)/48$  to obtain the lower bound on  $\delta t \geq 48\tilde{\Delta}/(\epsilon\beta f)$  that aligns with the bound in Corollary 6 (where  $\tilde{\Delta} = 2\Delta$ ), finally yielding the  $\tau_{\text{TG}}$  of the statement. Note that the coefficient tends to the value  $(1 - x)$  for increasing epoch lengths  $R \geq 144\tilde{\Delta}/(\epsilon\beta f)$ .  $\square$

Its main application is to the following corollary:

**Corollary 7.** *Consider the event that the execution of Ouroboros Chronos under the assumptions of Theorem 6 does not violate property CG with parameters  $\tau \in (0, 1]$ ,  $s \in \mathbb{N}$ . Let  $\tau_{\text{CG, glob}} \triangleq \tau \cdot \tau_{\text{TG}}$ . Consider a chain  $\mathcal{C}$  possessed by an alert party at a point in the execution where the party is at an onset of a (local) round and where the nominal time is  $t$ . Let further  $t_1, t_2$ , and  $\delta t$  be such that  $t_1 + \delta t \leq t_2 \leq t$ . Let  $s_{1_1}$  and  $s_{1_2}$  be the last slot numbers that  $P$  reported in the execution when nominal time was  $t_1$  (resp.  $t_2$ ) Then it must hold that  $|\mathcal{C}[s_{1_1} : s_{1_2}]| \geq \tau_{\text{CG, glob}} \cdot \delta t$  whenever  $\delta t \geq \max\{s/\tau, 48\tilde{\Delta}\}$ .*

*Proof (of Corollary 7).* By the previous Lemma, if the nominal time increases by  $\delta t$ , then in the view of alert party  $P$ , at least  $\tau_{\text{TG}} \cdot \delta t \geq s$  slots were reported (immediate by the property that all parties run at least at the baseline speed). Thus, by chain growth as of Definition 6.2, the increase in blocks between the reported logical slots  $s_{1_1}$  and  $s_{1_2}$  must be  $\tau_{\text{TG}} \cdot \delta t \cdot \tau = (\tau_{\text{TG}} \cdot \tau) \cdot \delta t$ .  $\square$

## L UC Realization

We are now ready to state the UC theorem in full detail, including all the parameters:

**Theorem 7.** *Let  $k$  be the common-prefix parameter,  $R$  the epoch-length parameter (constrained as required by Theorem 5) and  $\Delta = \Delta_{\text{net}} + \Delta_{\text{clock}}$ . Let  $\tau_{\text{CG}}$  be the chain growth coefficient as of Theorem 1, let  $\tau_{\text{CG, glob}}$  be the derived (nominal time) chain-growth coefficient as of Corollary 7, and let  $\mu$  be the chain quality coefficient as of Theorem 1. Under the constraints<sup>22</sup> of Theorem 6, the protocol Ouroboros Chronos realizes the ledger functionality, i.e., there exists a simulator that simulates the protocol execution in the ideal world perfectly except with negligible probability in the parameter  $k$  for  $R \geq \omega(\log k)$ , for the ledger parameters*

$$\begin{aligned} \text{windowSize} &= k; & \text{Delay} &= t_{\text{join}}; & \text{Delay}_{\text{tx}} &= 2\Delta; \\ \text{maxTime}_{\text{window}} &\geq \frac{\text{windowSize}}{\tau_{\text{CG}} \cdot \tau_{\text{CG, glob}}}; & \text{advBlcks}_{\text{window}} &\geq (1 - \mu)\text{windowSize}, \end{aligned}$$

and the clock-parameters

$$\begin{aligned} \text{shiftLB} &= -2\Delta; & \text{shiftUB} &= \Delta; & R_L &= R \\ \text{timeSlack}_{\text{total}} &= 2\Delta; & \text{timeSlack}_{\text{ep}} &= \Delta, \end{aligned}$$

and where the algorithms *Blockify*, *Validate*, and *predict-time* are instantiated as stated in Section M.2.

<sup>22</sup> Note that while we express the theorem as a constrained statement, it is possible to express the constraint as a (hybrid) functionality wrapper in the real world that enforces the constraints and leaves the environment unconstrained. For more details we refer to [5] and [4].

*Proof.* For the simulation part, we observe that compared to the simulator of the protocol Ouroboros Genesis in [4], the protocol Ouroboros Chronos does only introduce code such that the entire process of honest parties can still be simulated perfectly inside the simulator. All added procedures can be emulated by the simulator who is emulating the honest parties code, extracts their states and times, and sets the ledger parameter appropriately. We give the simulator in Section L.1.

We thus define the bad events that any constraints imposed by the above choice of parameters would prohibit the simulator in correctly setting the ledger state or the time (the events are called BAD-CP, BAD-CQ, BAD-CG, and BAD-TIME-RANGE in the simulation).

We first observe that violating the state parameters implies either violation of common-prefix, chain quality or chain-growth in the execution (i.e., we have a identical-until-bad simulation). Furthermore, the protocol is designed such that the activation pattern is still predictable by an efficiently computable predicate `predict-time`, since for the MAINTAIN-LEDGER it is by design fixed how many inputs are need to reach the `FinishRound` statement in the code.

Finally, the weak liveness of transactions holds since whenever a transaction is in the network at least  $\Delta < t_{\text{join}}$  rounds, it will eventually be included in the next high-quality block (i.e., a block with `hFlag = 1`) (and in the real-world by any alert party proposing a block) as long as the transaction is still valid. Considering that the analysis conditions no collisions among random oracle outputs, we obtain an upper bound of  $\exp(-\Omega(\kappa)) + \exp(\ln \text{poly}(\kappa) - \Omega(k)) + \exp(\ln \text{poly}(\kappa) - \Omega(R))$ , where  $\text{poly}(\kappa)$  denotes the polynomial upper bound on the runtime of  $\mathcal{Z}$  measured with respect to the security parameter  $\kappa$ . (Note that in particular, the parameters  $L$  and  $Q$  of the security bound can simply be upper bounded by this polynomial.) Combining this with the connection established in 6.7 settles that the chosen parameters do not impose a restriction on the ideal-world adversary except with negligible probability.

We next turn to the export-clock extension parameters. First, setting  $R_L = R$  is identical to the real world. Second, the simulator is given enough activation in every round s.t. whenever a synchronized party reaches a synchronization slot  $i \cdot R = i \cdot R_L$ , it can input a shift value. Next, by Lemma 5 it is clear that `shiftLB` =  $-2\Delta$  and `timeSlackep` =  $\Delta$  by Lemma 4. What remains to show is that we can essentially bound the (1) overall skew between two adjacent epochs by  $2\Delta$  and (2) that no party ever shifts its clock by more than  $\Delta$ . Both claims follow from directly from strengthening Lemma 5 as done in Lemma 9. By the preceding analysis, the probability that any constraint is violated, and thus BAD-TIME-RANGE is triggered, is also in this case bounded by a negligible function of the above form for our choice of parameters.  $\square$

**Lemma 9.** *Consider the same setting as in Lemma 5. Let `s1` be the synchronization slot of epoch `ep` and let  $\mathcal{S}_i^{\text{P}}$  be the set of beacons of an alert party  $\text{P}$  that is used for synchronization. Furthermore, assume that  $\text{Skew}_{\Delta}[\text{s1}]$  is not violated in this execution. Then it holds that*

1. *The shift `shift` party  $\text{P}$  computes is upper bounded by the maximal recommendation  $\text{recom}(\text{SB})$ ,  $\text{SB} \in \mathcal{S}_i^{\text{P}}$  for which  $\text{slotnum}(\text{SB}) = s$  for some  $s \in [(\text{ep} - 1)R, \dots, (\text{ep} - 1)R + R/6]$  and for which that it was created by an alert party  $\text{P}'$  in slot  $s$ . By the Lemma assumption, this is upper bounded by  $\Delta$ .*
2. *After the shift, party  $\text{P}$  reports a time-stamp that is at most  $2\Delta$  off of any alert party's reported time stamp.*

*Proof (Sketch).* The first part of the claim follows from the observation that `Timestamp(SB)` recorded by  $\text{P}$  is received after  $\text{P}'$  created the beacon. Consider the term  $\text{slotnum}(\text{SB}) - \text{Timestamp}(\text{SB})$  that this beacon contributes to the overall recommendation. If at the point the beacon was created the creator reported slot  $s$ , and party  $\text{P}'$  local timestamp (used to measure arrival times in epoch `ep`) differs by  $d$  to the creator's timestamp, then the recommendation is upper bounded by  $d$ , as delays can make the term only smaller and subsequent clock drifts do not influence the reported arrival times (as a party fetches once per observed round). The limited skew  $\text{Skew}_{\Delta}[s]$  at that slot  $s$  further says that  $d \leq \Delta$ . Since the alertly generated beacons are in the majority as argued in Lemma 5 the median is bounded by  $\Delta$  as well.

To prove the second item, note that it is sufficient to show that the difference in reported time stamp of party  $\text{P}$  in slot `s1` after the synchronization procedure to any alert party  $\text{P}'$  that has not yet made the clock adjustment for synchronization slot `s1` is at most  $2\Delta$  and cannot be further beyond that e.g., by doing some additional drift. First by 1., the party's shift is upper bounded by  $\Delta$ . Since by  $\text{Skew}_{\Delta}[\text{s1}]$ , any other alert party

$P'$  that has not yet passed synchronization slot  $\mathbf{s1}$ , will report a time stamp  $\mathbf{s1}' \geq \mathbf{s1} - \Delta$ . Finally, similarly to the maximal shift, the lower bound on the shift can be obtained by examining the recommendation computed by alertly generated beacons  $\text{slotnum}(\mathbf{SB}) - \text{Timestamp}(\mathbf{SB})$ . Analogous to the above case, if at the point of the execution the beacon was created the creator reported slot  $s$ , party  $P'$  local timestamp (used to measure arrival times in epoch  $\mathbf{ep}$ ) differs by at most  $d$  to the creator's timestamp, then the recommendation is lower bounded by  $d - \Delta$ . Similar to the arguments in Lemma 4, this worst case cannot be further amplified, because the drift allowed by  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  is an absolute additive constant over the entire execution (and therefore, either the true shift is in fact lower than the worst case and later maximized, or is maximal and cannot be further increased). Again, assuming  $\text{Skew}_{\Delta}[\mathbf{s1}]$  is not violated, this is lower bounded by  $-2\Delta$ . Since any alert party that has not yet made its adjustment reports a local time stamp larger equal to  $\mathbf{s1} - \Delta$ , the adjustment of  $P$  in this round will not increase the distance to any alert party that has not yet passed the synchronization slot to more than  $2\Delta$ . As above, this invariant remains true until everyone has made the adjustments (and then the stronger guarantees proven in previous sections apply).  $\square$

## L.1 Simulator

Below we present the simulator used in the proof that the UC implementation of Ouroboros Chronos securely realizes the ledger functionality  $\mathcal{G}_{\text{LEDGER}}$  with the clock extension. The simulator shares a lot of similarities with the simulator provided in [4] and is given below for the sake of concreteness.

### Simulator $\mathcal{S}_{\text{ledg}}$ (Part 1 - Main Structure)

#### Overview:

- The simulator internally emulates all local UC functionalities by running the code (and keeping the state) of  $\mathcal{F}_{\text{VRF}}$ ,  $\mathcal{F}_{\text{KES}}$ ,  $\mathcal{F}_{\text{INIT}}^{\Delta_{\text{net}}}$ ,  $\mathcal{F}_{\text{N-MC}}^{\text{bc}}$ ,  $\mathcal{F}_{\text{N-MC}}^{\text{tx}}$ , and  $\mathcal{F}_{\text{N-MC}}^{\text{sync}}$ , where  $\text{ref}_P := \langle v_P^{\text{vrf}}, v_P^{\text{kes}} \rangle$  identifies the address of party (i.e., ITI)  $P$  (and is given to the ledger when an ITI is registered).
- The simulator mimics the execution of Ouroboros Chronos for each honest party  $P$  (including their state and the interaction with the hybrids).
- The simulator emulates a view towards the adversary  $\mathcal{A}$  in a black-box way, i.e., by internally running adversary  $\mathcal{A}$  and simulating his interaction with the protocol (and hybrids) as detailed below for each hybrid. To simplify the description, we assume  $\mathcal{A}$  does not violate the theorem assumptions (as they are enforced by a wrapper  $\mathcal{W}_{\text{OG}}^{\text{PoS}}(\cdot)$  as in [4]).
- For global functionalities, the simulator simply relays the messages sent from  $\mathcal{A}$  to the global functionalities (and returns the generated replies). Recall that the ideal world consists of the dummy parties, the ledger functionality, the clock, and the random oracle.

#### Party sets:

- As defined in the main body of this paper, honest parties are categorized. We denote  $\mathcal{S}_{\text{alert}}$  the alert parties (synchronized and executing the protocol) and use  $\mathcal{S}_{\text{syncStalled}}$  shorthand for parties that are synchronized (and hence time aware and online) but stalled. Finally, we denote by  $\mathcal{P}_{DS}$  all honest but de-synchronized parties (both operational or stalled).
- For each registered honest party, the simulator maintains the local state containing in particular the local chain  $\mathcal{C}_{\text{loc}}^{(P)}$ , the time  $t_{\text{on}}$  it remembers when last being online. For each party  $P$ , the simulator maintains the objective (or nominal) round number  $\tau'_L := \tau_L + \delta_P$ , i.e., decomposed into the baseline time and the party-specific offset  $\delta_P$  bounded by  $\Delta_{\text{clock}}$ . For each party  $P$ , the simulator stores the reported time  $\mathbf{time}_P = (e, \text{localTime})$ , and the flags  $\text{updateState}_{P, \tau'_L}$ ,  $\text{updateTime}_{P, \tau'_L}$ , and  $\text{updateInitTime}_{P, \tau'_L}$  (initially false) to remember whether this party has completed its core maintenance tasks in its nominal round  $\tau'_L$  to update the state and its time (where the initial time for each party is a separate case), respectively. Note that an registered party is registered with all its local hybrids.
- Upon any activation, the simulator will query the current party set from the ledger, the clock, and the random oracle to evaluate in which category an honest party belongs to. If a new honest party is registered to the

ledger, it runs the initialization procedure for this party in each new round until the party is initialized ( $\mathbf{P.isInit}$  becomes true).

- We assume that the simulator queries upon any activation for the sequence  $\mathcal{I}_H^T$ , and the current (baseline) time  $\tau_L$  from the ledger. We note that the simulator is capable of determining  $\text{predict-time}(\cdot)$  of  $\mathcal{G}_{\text{LEDGER}}$  and hence the nominal time of the execution at any point. If the baseline advances, the simulator adjusts all offsets  $\delta_P$  accordingly (which it can do exactly as in  $\mathcal{G}_{\text{IMPERFLCLOCK}}$ ).

#### Messages involving the Clock:

- Upon receiving  $(\text{CLOCK-UPDATE}, \text{sid}_C, \mathbf{P})$  from  $\mathcal{G}_{\text{IMPERFLCLOCK}}$ , if  $\mathbf{P}$  is an honest registered party, then remember that this party has received such a clock update (and the environment gets an activation). Otherwise, send  $(\text{CLOCK-UPDATE}, \text{sid}_C, \mathbf{P})$  to  $\mathcal{A}$ .
- Upon receiving  $(\text{CLOCK-PUSH}, \text{sid}_C, \mathbf{P})$  from  $\mathcal{A}$ , remember that  $\mathbf{P}$  is allowed to locally advance in this round if allowed by  $\mathcal{G}_{\text{IMPERFLCLOCK}}$  (restricted by the maximal drift  $\Delta_{\text{clock}}$ ) and adjust  $\delta_P$ .

#### Messages from the Ledger:

- Upon receiving  $(\text{Respond}, (\text{start}, \text{sid}))$  from  $\mathcal{G}_{\text{LEDGER}}$ , send  $(\text{Respond}, (\text{DefineOffset}, \text{sid}))$  in the name of  $\mathcal{F}_{\text{INIT}}^\Delta$  to the adversary. Upon receiving the response  $(\text{DefineOffset}, \text{sid}, o_1, \dots, o_n)$ ,  $o_i \in [0, \dots, \Delta]$ , store the values and relay the answer to the simulated instance of  $\mathcal{F}_{\text{INIT}}^\Delta$ .
- Upon receiving  $(\text{SUBMIT}, \text{BTX})$  from  $\mathcal{G}_{\text{LEDGER}}$  where  $\text{BTX} := (\text{tx}, \text{txid}, \tau, \mathbf{P})$  forward  $(\text{MULTICAST}, \text{sid}, \text{tx})$  to the simulated network  $\mathcal{F}_{\text{N-MC}}^{\text{tx}}$  in the name of  $\mathbf{P}$ . Output the answer of  $\mathcal{F}_{\text{N-MC}}$  to the adversary.
- Upon receiving  $(\text{MAINTAIN-LEDGER}, \text{sid}, \text{minerID})$  from  $\mathcal{G}_{\text{LEDGER}}$ , extract from  $\mathcal{I}_H^T$  the party  $\mathbf{P}$  that issued this query. If  $\mathbf{P}$  has already completed its round-task, then ignore this request. Otherwise, execute  $\text{SIMULATEMAINTENANCE}(\mathbf{P}, \tau_L)$ .

### Simulator $\mathcal{S}_{\text{ledg}}$ (Part 2 - Black-Box Interaction)

*Simulation of Functionality  $\mathcal{F}_{\text{INIT}}$  towards  $\mathcal{A}$ :*

- The simulator relays back and forth the communication between the (internally emulated)  $\mathcal{F}_{\text{INIT}}^{\Delta_{\text{net}}}$  functionality and the adversary  $\mathcal{A}$  acting on behalf of a corrupted party.
- If at time  $\tau_L = 0$ , a corrupted party  $\mathbf{P} \in \mathcal{S}_{\text{initStake}}$  registers via  $(\text{ver\_keys}, \text{sid}, \mathbf{P}, v_P^{\text{vrf}}, v_P^{\text{kes}})$  to  $\mathcal{F}_{\text{INIT}}$ , then input  $(\text{REGISTER}, \text{sid})$  to  $\mathcal{G}_{\text{LEDGER}}$  on behalf of  $\mathbf{P}$ .

*Simulation of the Functionalities  $\mathcal{F}_{\text{KES}}$  and  $\mathcal{F}_{\text{VRF}}$  towards  $\mathcal{A}$ :*

- The simulator relays back and forth the communication between the (internally emulated) hybrids and the adversary  $\mathcal{A}$  (either direct communication, communication to  $\mathcal{A}$  caused by emulating the actions of honest parties, or communication of  $\mathcal{A}$  on behalf of a corrupted party).

*Simulation of the Network  $\mathcal{F}_{\text{N-MC}}^{\text{bc}}$  (over which chains are sent) towards  $\mathcal{A}$ :*

- Upon receiving  $(\text{MULTICAST}, \text{sid}, (\mathcal{C}_{i_1}, U_{i_1}), \dots, (\mathcal{C}_{i_\ell}, U_{i_\ell}))$  with a list of chains and corresponding parties from  $\mathcal{A}$  (or on behalf some *corrupted*  $P \in \mathcal{P}_{\text{net}}$ ), then do the following:
  1. Relay this input to the simulate network functionality and record its response to  $\mathcal{A}$ .
  2. Execute  $\text{EXTENDLEDGERSTATE}(\tau_L)$
  3. Provide  $\mathcal{A}$  with the recorded output of the simulated network.
- Upon receiving  $(\text{MULTICAST}, \text{sid}, \mathcal{C})$  from  $\mathcal{A}$  on behalf of some *corrupted* party  $P$ , then do the following:
  1. Relay this input to the simulate network functionality and record its response to  $\mathcal{A}$ .
  2. Execute  $\text{EXTENDLEDGERSTATE}(\tau_L)$
  3. Provide  $\mathcal{A}$  with the recorded output of the simulated network.
- Upon receiving  $(\text{FETCH}, \text{sid})$  from  $\mathcal{A}$  on behalf some *corrupted*  $P \in \mathcal{P}_{\text{net}}$  forward the request to the simulated  $\mathcal{F}_{\text{N-MC}}^{\text{bc}}$  and return whatever is returned to  $\mathcal{A}$ .
- Upon receiving  $(\text{DELAYS}, \text{sid}, (T_{\text{mid}_{i_1}}, \text{mid}_{i_1}), \dots, (T_{\text{mid}_{i_\ell}}, \text{mid}_{i_\ell}))$  from  $\mathcal{A}$ : Forward the request to the simulated  $\mathcal{F}_{\text{N-MC}}^{\text{bc}}$  and record the answer to  $\mathcal{A}$ . Before giving this answer to  $\mathcal{A}$ , query the ledger state  $\text{state}$  and execute  $\text{ADJUSTVIEW}(\text{state}, \tau_L)$ .

- Upon receiving (SWAP, sid, mid, mid') from  $\mathcal{A}$ : Forward the request to the simulated  $\mathcal{F}_{N-MC}^{bc}$  and record the answer to  $\mathcal{A}$ . Before giving this answer to  $\mathcal{A}$ , query the ledger state **state** and execute  $\text{ADJUSTVIEW}(\mathbf{state}, \tau_L)$ .

*Simulation of the Network  $\mathcal{F}_{N-MC}^{tx}$  (over which transactions are sent) towards  $\mathcal{A}$ :*

- Upon receiving (MULTICAST, sid,  $m$ ) from  $\mathcal{A}$  with list a transaction  $m$  from  $\mathcal{A}$  on behalf some *corrupted*  $P \in \mathcal{P}_{net}$ , then do the following:
  1. Submit the transaction to the ledger on behalf of this corrupted party, and receive for the transaction id txid.
  2. Forward the request to the internally simulated  $\mathcal{F}_{N-MC}^{tx}$ , which replies for each message with a message-ID mid
  3. Remember the association between mid and the corresponding txid
  4. Provide  $\mathcal{A}$  with whatever the network outputs.
- Upon receiving (FETCH, sid) from  $\mathcal{A}$  on behalf some *corrupted*  $P \in \mathcal{P}_{net}$  forward the request to the simulated  $\mathcal{F}_{N-MC}^{tx}$  and return whatever is returned to  $\mathcal{A}$ .
- Upon receiving (DELAYS, sid,  $(T_{mid_{i_1}}, mid_{i_1}), \dots, (T_{mid_{i_\ell}}, mid_{i_\ell})$ ) from  $\mathcal{A}$  forward the request to the simulated  $\mathcal{F}_{N-MC}^{tx}$  and return whatever is returned to  $\mathcal{A}$ .
- Upon receiving (SWAP, sid, mid, mid') from  $\mathcal{A}$  forward the request to the simulated  $\mathcal{F}_{N-MC}^{tx}$  and return whatever is returned to  $\mathcal{A}$ .

*Simulation of the Network  $\mathcal{F}_{N-MC}^{sync}$  (over which beacons are sent) towards  $\mathcal{A}$ :*

- Upon receiving (MULTICAST, sid,  $m$ ) from  $\mathcal{A}$  with a beacon  $m$  from  $\mathcal{A}$  on behalf some *corrupted*  $P \in \mathcal{P}_{net}$ , then do the following:
  1. Forward the request to the internally simulated  $\mathcal{F}_{N-MC}^{sync}$ , which replies for each message with a message-ID mid
  2. Remember the association between each mid and the corresponding beacon.
  3. Provide  $\mathcal{A}$  with whatever the network outputs.
- Upon receiving (FETCH, sid) from  $\mathcal{A}$  on behalf some *corrupted*  $P \in \mathcal{P}_{net}$  behave analogously to above for  $\mathcal{F}_{N-MC}^{tx}$ .
- Upon receiving (DELAYS, sid,  $(T_{mid_{i_1}}, mid_{i_1}), \dots, (T_{mid_{i_\ell}}, mid_{i_\ell})$ ) from  $\mathcal{A}$  behave analogously to above for  $\mathcal{F}_{N-MC}^{tx}$ .
- Upon receiving (SWAP, sid, mid, mid') from  $\mathcal{A}$  behave analogously to above for  $\mathcal{F}_{N-MC}^{tx}$ .

### Simulator $\mathcal{S}_{ledg}$ (Part 3 - Internal Procedures)

**procedure** SIMULATEMAINTENANCE( $P, \tau_L$ )

Simulate the (in the UC interruptible manner) the maintenance procedure of party  $P$  as in the protocol in round  $\tau'_L = \tau_L + \delta_P$  when the party reports localtime  $P.localTime$ , i.e., run  $\text{LedgerMaintenance}(\mathcal{C}_{loc}, P, sid, k, s, R, f)$  for this simulated party.

**if** party  $P$  gives up activation **then then**

**if** party  $P$  has completed  $\text{JoinProc}(\cdot)$  and  $\text{updateInitTime}_{P, \tau'_L}$  is false **then**

    Execute  $\text{ADJUSTTIME}(\tau_L)$  and then set  $\text{updateInitTime}_{P, \tau'_L} \leftarrow \text{true}$ .

**end if**

**if** party  $P$  has reached the instruction  $\text{SelectChain}(\cdot)$  and  $\text{updateState}_{P, \tau'_L}$  is false **then**

    Execute  $\text{EXTENDLEDGERSTATE}(\tau_L)$  and then set  $\text{updateState}_{P, \tau'_L} \leftarrow \text{true}$ .

**end if**

**if** party  $P$  has reached the instruction  $\text{SyncProc}(\cdot)$  and  $\text{updateTime}_{P, \tau'_L}$  is false **then**

    Execute  $\text{ADJUSTTIME}(\tau_L)$  and then set  $\text{updateTime}_{P, \tau'_L} \leftarrow \text{true}$ .

**end if**

**if** party  $P$  has reached the instruction  $\text{FinishRound}(P)$  in round  $\tau'_L$  **then**

    Send  $(\text{CLOCK-UPDATE}, sid_C, P)$  to  $\mathcal{A}$  if  $\mathcal{S}_{ledg}$  has received such an input in round  $\tau'_L$

**end if**

  Return activation to  $\mathcal{A}$

**end if**



**end procedure**

**procedure** EXTENDLEDGERSTATE( $\tau_L$ )

**for** each synchronized party  $P \in \mathcal{S}_{\text{alert}} \cup \mathcal{S}_{\text{syncStalled}}$  of round `localTime` **do**

Let  $\mathcal{C}_{\text{loc}}^{(P)}$  be the party's currently stored local chain.

// Note: In the following the internally simulated party state is not changed

Determine the number of fetches  $\rho^{(P)} \in \{0, 1\}$  this party is still going to make in this round  $\tau'_L := \tau_L + \delta_P$ .

If  $\rho^{(P)} > 0$  then let  $\mathcal{C}_1^{(P)}, \dots, \mathcal{C}_k^{(P)}$  be the chains contained in the receiver buffer  $\mathcal{M}^{(P)}$  of  $\mathcal{F}_{\text{N-MC}}^{\text{bc}}$  with delay at most  $\rho^{(P)}$ .

Re-evaluate  $\mathcal{C}_P \leftarrow \text{SelectChain}()$  using the additional chains as well and let this resulting chain's encoded state be `stP`.

**end for**

Let `st` be the longest state among all such states `stP`,  $P \in \mathcal{S}_{\text{alert}} \cup \mathcal{S}_{\text{syncStalled}}$  from above.

Compare `st`<sup>[k</sup> with the current state `state` of the ledger

**if** `|state| > |st`<sup>[k</sup> **then** // Only pointers need adjustments

Execute `ADJUSTVIEW(state)`

**end if**

**if** `state` is not a prefix of `st`<sup>[k</sup> **then** // Simulation fails

**Abort** simulation: consistency violation among synchronized parties. // Event `BAD-CPk`

**end if**

Define the difference `diff` to be the block sequence s.t. `state||diff = st`<sup>[k</sup>.

Parse `diff := diff1} \dots ||diffn`.

**for**  $j = 1$  to  $n$  **do**

Map each transaction `tx` in this block to its unique transaction ID `txid`. If a transaction does not yet have a `txid`, then submit it to the ledger first and receive the corresponding `txid` from  $\mathcal{G}_{\text{LEDGER}}$

Let `listj = (txidj,1}, \dots, txidj,\ell_j})` be the corresponding list for this block `diffj`

**if** coinbase `txidj,1}` specifies a party honest at block creation time **then**

`hFlagj ← 1`

**else**

`hFlagj ← 0`

**end if**

Output `(NEXT-BLOCK, hFlagj}, listj})` to  $\mathcal{G}_{\text{LEDGER}}$  (receiving `(NEXT-BLOCK, ok)` as an immediate answer)

**end for**

**if** Fraction of blocks with `hFlag = 0` in the recent  $k$  blocks  $> 1 - \mu$  **then**

**Abort** simulation: chain quality violation. // Event `BAD-CQ\mu,k}`

**else if** State increases less than  $k$  blocks during the last  $\frac{k}{\tau_{\text{CG}}}$  rounds **then**

**Abort** simulation: chain growth violation. // Event `BAD-CG\tau_{\text{CG}},k/\tau_{\text{CG}}`

**end if**

// If no bad event occurs, we can adjust pointers into this new state.

Execute `ADJUSTVIEW(state||diff)`

**end procedure**

**procedure** ADJUSTTIME( $P, \tau_L$ )

Let  $\tau'_L := \tau_L + \delta_P$  be the current (objective) round of this party.

**if**  $P$  completed `JoinProc` in this round  $\tau'_L$  **then**

// Note that this party is about to become synchronized and to report time.

Take the simulated timestamp `timeP` and send `(SET-TIME, sid, P, timeP})` to  $\mathcal{G}_{\text{LEDGER}}$ .

**end if**

**if**  $\tau'_L = 0$  and  $P$  is an initial stakeholder  $U_i$  **then**

Send `(APPLY-SHIFT, sid, (Ui}, -oi})` to  $\mathcal{G}_{\text{LEDGER}}$ , where the  $o_i$  are the initial offsets by  $\mathcal{A}$  as recorded above.

**else if**  $P.\text{localTime} \bmod R = 0$  **then**

Take the simulated timestamp `timeP` and the simulated shift `shift` of this party.

**if** The range of timestamps of parties in  $\mathcal{S}_{\text{alert}} \cup \mathcal{S}_{\text{syncStalled}}$  is invalid **then**

```

    Abort simulation: time-range violation. // Event BAD-TIME-RANGE
else
    Send (APPLY-SHIFT, sid, (P, shift)) to  $\mathcal{G}_{\text{LEDGER}}$ .
end if
end if
end procedure

```

## M Glossary: Protocol Variables and Ledger Parameters

### M.1 Main State Variables of Ouroboros Chronos

The variables defining the state of a protocol participant are summarized below.

Variable	Description
<code>localTime, s1</code>	The party's current time-stamp. In the staking context we call the time slot.
<code>ep</code>	The epoch that <code>s1</code> belongs to.
<code>C<sub>loc</sub></code>	The local chain the party adopts based on which it does staking and exports the ledger state.
<code>isInit</code>	A variable to keep track of whether initialization is complete.
<code>t<sub>work</sub></code>	A value to steer when the party executes the staking procedure for the next time.
<code>buffer</code>	The buffer of transactions.
<code>futureChains</code>	A buffer to store chains that are not yet processed, for example because they contain blocks that belong to the logical future of this party.
<code>Timestamp<sub>SB</sub>(·)</code>	A map that assigns to each synchronization beacon a pair $(a, b)$ , where $a$ is a numerical value (the arrival time) and $b$ is an indication of whether $a$ is final or not.
<code>Timestamp(·)</code>	Shorthand for the first (and numerical) element of the pair <code>Timestamp<sub>SB</sub>(·)</code> .
<code>lastTick</code>	The last tick received from $\mathcal{G}_{\text{PERFLCLOCK}}$ . Used to infer when a round change occurs.
<code>isSync</code>	A party stores its synchronization status, as it can infer when its time and state become reliable.
<code>EpochUpdate(·)</code>	An function table to remember which clock adjustments have been done already. Used to update beacon arrival times.
<code>fetchCompleted</code>	A variable to store whether the round messages have been fetched.
<code>lastTimeAlert</code>	The local time stamp the party was alert the last time. Used for rejoining if the party was only stalled.
$T_p^{\text{ep}}, T_p^{\text{ep.bc}}$	The thresholds of this party to evaluate slot leadership (and beacon production and validity) in (current) epoch <code>ep</code> .
$v_p^{\text{vrf}}, v_p^{\text{kes}}$	The public keys of this party to interact with $\mathcal{F}_{\text{KES}}$ and $\mathcal{F}_{\text{VRF}}$ .

**Fig. 4.** Overview of the main state variables of Ouroboros Chronos.

## M.2 Main Ledger Elements

The state variables of the new ledger functionality are given below.

Core Ledger Parameter	Description
<code>windowSize</code>	The window size (number of blocks) of the sliding window. In the realization statement, it is typically set to the common-prefix parameter.
<code>Validate</code>	Decides on the validity of a transaction with respect to the current state. Used to clean the buffer of transactions. If the protocol fixes a validation predicate, say <code>ValidTx<sub>OC</sub></code> , then the realization statement holds with <code>Validate(BTX, state, buffer) := ValidTx<sub>OC</sub>(tx, state)</code> .
<code>Blockify</code>	The function to format the ledger state output. If the protocol fixes a particular function, say <code>blockify<sub>OC</sub></code> , the ledger will use the same in the realization proof.
<code>predict-time</code>	The function to predict the real-world time advancement. Ouroboros Chronos has a predictable time-advancement <code>predict-time<sub>OC</sub></code> as it can be inferred by design when the protocol will call <code>FinishRound</code> in each round, when given a fixed number of activations that depends on the local time-stamp of this party.
<code>Delay</code>	A general delay parameter for the time it takes for a newly joining (after the onset of the computation) miner to become synchronized. In this paper, it corresponds to the duration of the joining procedure.
Policy Parameter (ExtendPolicy)	Description
<code>maxTime<sub>window</sub></code>	Minimal Growth: In <code>maxTime<sub>window</sub></code> rounds at least <code>windowSize</code> blocks have to be inserted into the ledger state. The value in the realization proof will depend on the chain-growth property.
<code>advBlcks<sub>window</sub></code>	A limit <code>advBlcks<sub>window</sub></code> of adversarial blocks (i.e., contributed blocks that do not need to employ higher standards) in each window of <code>windowSize</code> state blocks. This ensures a minimal fraction of blocks that contain all old and valid transactions. The value in the realization proof will depend on the chain-quality property.
<code>Delay<sub>tx</sub></code>	An extra parameter to define when a transaction is old. In this work, this will be much less than <code>Delay</code> as it will only depend on the network delay.
Export-Time Parameter	Description
<code>time<sub>P</sub></code>	A variable that will represent the (idealized) clock value that the party reports as its local time. A party will export pairs $(e, t)$ , where $t$ is the current local time, and $e$ is the epoch.
<code>shiftLB, shiftUB</code>	Limits on the shift values an adversary can impose at epoch boundaries
$R_L$	The parameter characterizing epoch boundaries: if a party's time-stamp $(e, t)$ is such that $t = iR_L + 1$ (for the first time), then the party moves to the next epoch.
<code>timeSlack<sub>total</sub></code>	An upper bound between $t$ and $t'$ of two synchronized parties $P$ and $P'$ reporting $(e, t)$ and $(e', t')$ as their respective time <code>time<sub>P</sub></code> and <code>time<sub>P'</sub></code> , respectively.
<code>timeSlack<sub>ep</sub></code>	An upper bound between $t$ and $t'$ of two synchronized parties $P$ and $P'$ reporting $(e, t)$ and $(e', t')$ as their respective time <code>time<sub>P</sub></code> and <code>time<sub>P'</sub></code> , respectively whenever $e = e'$

**Fig. 5.** Overview of main ledger elements such as parameters and state variables. As in [5, Definition 2], we always assume that `blockifyOC` and `ValidTxOC` do not disqualify each other.