Self-Stabilising Byzantine Clock Synchronisation Is Almost as Easy as Consensus

CHRISTOPH LENZEN, Max Planck Institute for Informatics, Saarland Informatics Campus, Germany
JOEL RYBICKI, Institute of Science and Technology Austria (IST Austria), Austria

We give fault-tolerant algorithms for establishing synchrony in distributed systems in which each of the n nodes has its own clock. Our algorithms operate in a very strong fault model: we require self-stabilisation, i.e., the initial state of the system may be arbitrary, and there can be up to $f < n/3$ ongoing Byzantine faults, i.e., nodes that deviate from the protocol in an arbitrary manner. Furthermore, we assume that the local clocks of the nodes may progress at different speeds (clock drift) and communication has bounded delay. In this model, we study the pulse synchronisation problem, where the task is to guarantee that eventually all correct nodes generate well-separated local pulse events (i.e., unlabelled logical clock ticks) in a synchronised manner.

Compared to prior work, we achieve *exponential* improvements in stabilisation time and the number of communicated bits, and give the first sublinear-time algorithm for the problem:

- In the deterministic setting, the state-of-the-art solutions stabilise in time $\Theta(f)$ and have each node broadcast $\Theta(f \log f)$ bits per time unit. We exponentially reduce the number of bits broadcasted per time unit to $\Theta(\log f)$ while retaining the same stabilisation time.
- In the randomised setting, the state-of-the-art solutions stabilise in time $\Theta(f)$ and have each node broadcast $O(1)$ bits per time unit. We exponentially reduce the stabilisation time to $\text{polylog } f$ while each node broadcasts $\text{polylog } f$ bits per time unit.

These results are obtained by means of a recursive approach reducing the above task of *self-stabilising* pulse synchronisation in the *bounded-delay* model to *non-self-stabilising* binary consensus in the *synchronous* model. In general, our approach introduces at most logarithmic overheads in terms of stabilisation time and broadcasted bits over the underlying consensus routine.

CCS Concepts: • **Theory of computation** → **Distributed algorithms**;

Additional Key Words and Phrases: Transient faults, byzantine faults, clock drift, agreement

ACM Reference format:

Christoph Lenzen and Joel Rybicki. 2019. Self-Stabilising Byzantine Clock Synchronisation Is Almost as Easy as Consensus. *J. ACM* 66, 5, Article 32 (August 2019), 56 pages.
<https://doi.org/10.1145/3339471>

This manuscript is an extended and revised version of a conference report that appeared in Proceedings of the 31st International Symposium on Distributed Computing (DISC 2017).

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 716562) and funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 754411.

Authors' addresses: C. Lenzen, Department of Algorithms and Complexity, Max Planck Institute for Informatics, Saarland Informatics Campus, Campus E1 4, Saarbrücken, 66123, Germany; email: clenzen@mpi-inf.mpg.de; J. Rybicki, Institute of Science and Technology Austria (IST Austria), Am Campus 1, Klosterneuburg, 3400, Austria; email: joel.rybicki@ist.ac.at.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2019 Copyright held by the owner/author(s).

0004-5411/2019/08-ART32

<https://doi.org/10.1145/3339471>

1 INTRODUCTION

Many of the most fundamental problems in distributed computing relate to timing and fault tolerance. Even though most distributed systems are inherently asynchronous, it is often convenient to design such systems by assuming some degree of synchrony provided by reliable global or distributed clocks. For example, the vast majority of existing Very Large Scale Integrated (VLSI) circuits operate according to the synchronous paradigm: an internal clock signal is distributed throughout the chip neatly controlling alternation between computation and communication steps. Of course, establishing the synchronous abstraction is of high interest in numerous other large-scale distributed systems, as it makes the design of algorithms considerably easier.

However, as the accuracy and availability of the clock signal is typically one of the most basic assumptions, clocking errors affect system behavior in unpredictable ways that are often hard—if not impossible—to tackle at higher system layers. Therefore, *reliably* generating and distributing a joint clock is an essential task in distributed systems. Unfortunately, the cost of providing fault-tolerant synchronisation and clocking is still poorly understood.

1.1 Pulse Synchronisation

In this work, we study the *self-stabilising Byzantine pulse synchronisation* problem [13, 16], which requires the system to achieve synchronisation despite severe faults. We assume a fully connected message-passing system of n nodes, where

- (1) an unbounded number of transient faults may occur anywhere in the network, and
- (2) up to $f < n/3$ of the nodes can be faulty and exhibit *arbitrary* ongoing misbehaviour.

In particular, the transient faults may arbitrarily corrupt the state of the nodes and result in loss of synchrony. Moreover, the nodes that remain faulty may deviate from any given protocol, behave adversarially, and collude to disrupt the other nodes by sending them *different* misinformation even after transient faults have ceased. Note that this also covers faults of the communication network, as we may map faults of communication links to one of their respective endpoints. The goal is now to (re-)establish synchronisation once transient faults cease, despite up to $f < n/3$ Byzantine nodes. That is, we need to consider algorithms that are simultaneously (1) self-stabilising [7, 15] and (2) Byzantine fault-tolerant [23].

More specifically, the problem is as follows: after transient faults cease, no matter what is the initial state of the system, the choice of up to $f < n/3$ faulty nodes, and the behaviour of the faulty nodes, we require that after a bounded *stabilisation time* all the *non-faulty* nodes must generate pulses that

- occur almost simultaneously at each correctly operating node (i.e., have small *skew*), and
- satisfy specified minimum and maximum frequency bounds (*accuracy*).

While the system may have arbitrary behaviour during the initial stabilisation phase due to the effects of transient faults, eventually the above conditions provide synchronised unlabelled clock ticks for all non-faulty nodes as shown in Figure 1.

In order to meet these requirements, it is necessary that nodes can estimate the progress of time. To this end, we assume that nodes are equipped with (continuous, real-valued) hardware clocks that run at speeds that may vary arbitrarily within 1 and ϑ , where $\vartheta \in O(1)$. That is, we normalise minimum clock speed to 1 and assume that the clocks have drift bounded by a constant. Observe that in an asynchronous system, i.e., one in which communication and/or computation may take unknown and unbounded time, even perfect clocks are insufficient to ensure any relative timing guarantees between the actions of different nodes. Therefore, we additionally assume that the nodes can send messages to each other that are received and processed within at most $d \in \Theta(1)$

Fig. 1. An example execution of a pulse synchronisation algorithm. After the initial stabilisation phase (delimited by the vertical lines), correct nodes start generating well-separated pulses with bounded skew and accuracy.

time. The clock speeds and message delays can behave adversarially within the respective bounds given by ϑ and d .

In summary, this yields a highly adversarial model of computing, where further restrictions would render the task infeasible:

- (1) transient faults are arbitrary and may involve the entire network,
- (2) ongoing faults are arbitrary, cover erroneous behavior of both the nodes and the communication links, and the problem is not solvable if $f \geq n/3$ [11], and
- (3) without any bounds on the accuracy of local clocks and on the communication delay, good synchronisation cannot be achieved: Even without clock drift, unbounded message delays lead to unbounded skew [27], and if clocks have unbounded drift, trivial indistinguishability arguments show that no bounds on pulse frequency can be maintained.

1.2 Background and Related Work

If one takes any one of the elements described above out of the picture, then this greatly simplifies the problem. Without permanent/ongoing faults, the problem becomes trivial: it suffices to have all nodes follow a designated leader. Without transient faults [22], straightforward solutions are given by elegant classics [31, 32], where [32] also guarantees asymptotically optimal skew [27]. Taking the uncertainty of unknown message delays and drifting clocks out of the equation leads to the so-called digital clock synchronisation problem [3, 14, 24, 26], where communication proceeds in synchronous rounds and the task is to agree on a consistent (bounded) round counter. While this abstraction is unrealistic as a basic system model, it yields conceptual insights into the pulse synchronisation problem in the bounded-delay model. Moreover, it is useful to assign numbers to pulses after pulse synchronisation is solved, in order to get a fully-fledged shared system-wide clock [25].

In contrast to these relaxed problem formulations, the pulse synchronisation problem was initially considered to be very challenging—if not impossible—to solve. In a seminal article, Dolev and Welch [16] proved otherwise, albeit with an algorithm having an impractical exponential stabilisation time. In a subsequent line of work, the stabilisation time was reduced to polynomial [6] and then linear in f [9]. However, the linear-time algorithm relies on simulating multiple instances of synchronous *consensus* algorithms [28] concurrently, which results in a high communication complexity.

The consensus problem [23, 28] is one of the fundamental primitives in fault-tolerant computing. Most relevant to this work is synchronous binary consensus with (up to f) Byzantine faults. Here, node v is given an input $x(v) \in \{0, 1\}$, and it must output $y(v) \in \{0, 1\}$ such that the following properties hold:

- (1) *Agreement.* There exists $y \in \{0, 1\}$ such that $y(v) = y$ for all correct nodes v .
- (2) *Validity.* If for $x \in \{0, 1\}$ it holds that $x(v) = x$ for all correct nodes v , then $y = x$.
- (3) *Termination.* All correct nodes eventually decide on $y(v)$ and terminate.

In this setting, two of the above main obstacles are not present: the system is properly initialised (no self-stabilisation required) and computation proceeds in synchronous rounds, i.e., well-ordered compute-send-receive cycles. This confines the task to understanding how to deal with the interference from Byzantine nodes. Synchronous consensus is extremely well studied; see, e.g., [30] for a survey. It is known that precisely $\lfloor (n-1)/3 \rfloor$ faults can be tolerated in a system of n nodes [28], $\Omega(nf)$ messages need to be sent in total [10], the connectivity of the communication network must be at least $2f + 1$ [8], deterministic algorithms require $f + 1$ rounds [1, 19], and randomised algorithms can solve the problem in constant expected time [18]. In contrast, no non-trivial lower bounds on the time or communication complexity of pulse synchronisation are known.

The linear-time pulse synchronisation algorithm in [9] relies on simulating (up to) one synchronous consensus instance for each node simultaneously. Accordingly, this protocol requires each node to broadcast $\Theta(f \log f)$ bits per time unit. Moreover, the use of *deterministic* consensus is crucial, as failure of any consensus instance to generate correct output within a prespecified time bound may result in loss of synchrony, i.e., the algorithm would fail *after* apparent stabilisation. In [13], these obstacles were overcome by avoiding the use of consensus by reducing the pulse synchronisation problem to the easier task of generating at least one well-separated “resynchronisation point”, which is roughly uniformly distributed within any period of $\Theta(f)$ time. This can be achieved by trying to initiate such a resynchronisation point at random times, in combination with threshold voting and locally checked timing constraints to rein in the influence of Byzantine nodes. In a way, this seems much simpler than solving consensus, but the randomisation used to obtain a suitable resynchronisation point strongly reminds of the power provided by shared coins [2, 3, 18, 29]—and this is exactly what the core routine of the expected constant-round consensus algorithm from [18] provides.

1.3 Contributions

Our main result is a framework that reduces pulse synchronisation to an arbitrary synchronous binary consensus routine at very small overheads. In other words, given *any* efficient algorithm that solves consensus in the standard synchronous model of computing *without* self-stabilisation, we show how to obtain an efficient algorithm that solves the *self-stabilising* pulse synchronisation problem in the bounded-delay model with clock drift.

While we build upon existing techniques, our approach has many key differences. First of all, while Dolev et al. [13] also utilise the concept of resynchronisation pulses, these are generated probabilistically. Moreover, their approach has an inherent time bound of $\Omega(f)$ for generating such pulses. In contrast, we devise a new recursive scheme that allows us to (1) *deterministically* generate resynchronisation pulses in $\Theta(f)$ time and (2) *probabilistically* generate resynchronisation pulses in $o(f)$ time. To construct algorithms that generate resynchronisation pulses, we employ resilience boosting and filtering techniques inspired by our recent line of work on digital clock synchronisation in the *synchronous* model [24, 26]. One of its main motivations was to gain a better understanding of the linear time/communication complexity barrier that research on pulse synchronisation ran into, without being distracted by the additional timing uncertainties due to communication delay and clock drift. The challenge here is to port these newly developed tools from the synchronous model to the bounded-delay bounded-drift model in a way that keeps them in working condition.

Table 1. Summary of Pulse Synchronisation Algorithms for $f \in \Theta(n)$

time	bits	type	notes	reference
$\text{poly } f$	$O(\log f)$	det.		[6]
$O(f)$	$O(f \log f)$	det.		[9]
$O(f)$	$O(\log f)$	det.		this work and [4]
$2^{O(f)}$	$O(1)$	rand.	adversary cannot predict coin flips	[16]
$O(f)$	$O(1)$	rand.	adversary cannot predict coin flips	[13]
$\text{polylog } f$	$\text{polylog } f$	rand.	private channels, (*)	this work and [21]
$O(\log f)$	$\text{poly } f$	rand.	private channels	this work and [18]

For each respective algorithm, the first two columns give the stabilisation time and the number of bits broadcasted by a node per time unit. The third column denotes whether algorithm is deterministic or randomised. The randomised algorithms stabilise in the given time with high probability. The fourth column indicates additional details or model assumptions. All algorithms tolerate $f < n/3$ faulty nodes except for (*), where it is required that $f < n/(3 + \varepsilon)$ for an arbitrary, but fixed constant $\varepsilon > 0$.

The key to efficiency is a recursive approach, where each node participates in only $\lceil \log f \rceil$ consensus instances, one for each level of recursion. On each level, the overhead of the reduction over a call to the consensus routine is a constant multiplicative factor both in time and bit complexity; concretely, this means that both complexities increase by overall factors of $O(\log f)$. Applying suitable consensus routines yields *exponential improvements* in bit complexity of deterministic and time complexity of randomised solutions, respectively:

- (1) In the deterministic setting, we exponentially reduce the number of bits each node broadcasts per time unit to $\Theta(\log f)$, while retaining $\Theta(f)$ stabilisation time. This is achieved by employing the phase king algorithm [4] in our construction.
- (2) In the randomised setting, we exponentially reduce the stabilisation time to $\text{polylog } f$, where each node broadcasts $\text{polylog } f$ bits per time unit. This is achieved using the algorithm by King and Saia [21]. We note that this slightly reduces resilience to $f < n/(3 + \varepsilon)$ for any fixed constant $\varepsilon > 0$ and requires private communication channels.
- (3) In the randomised setting, we can also obtain a stabilisation time of $O(\log f)$, polynomial communication complexity, and optimal resilience of $f < n/3$ by assuming private communication channels. This is achieved using the consensus routine of Feldman and Micali [18]. This almost settles the open question by Ben-Or et al. [3] whether pulse synchronisation can be solved in expected constant time.

The running time bounds of the randomised algorithms (2) and (3) hold with high probability and the additional assumptions on resilience and private communication channels are inherited from the employed consensus routines. Here, private communication channels mean that Byzantine nodes must make their decision on which messages to send in round r based on knowledge of the algorithm, inputs, and all messages faulty nodes receive up to and including round r . The probability distribution is then over the independent internal randomness of the correct nodes (which the adversary can only observe indirectly) and any possible randomness of the adversary. Our framework does not impose these additional assumptions: stabilisation is guaranteed for $f < n/3$ on each recursive level of our framework as soon as the underlying consensus routine succeeds (within prespecified time bounds) constantly many times in a row. Our results and prior work are summarised in Table 1.

Regardless of the employed consensus routine, we achieve a skew of $2d$, where d is the maximum message delay. This is optimal in our model, but overly pessimistic if the sum of communication and computation delay is not between 0 and d , but from (d^-, d^+) , where $d^+ - d^- \ll d^+$. In terms

of d^+ and d^- , a skew of $\Theta(d^+ - d^-)$ is asymptotically optimal [27, 32]. We remark that in [20], it is shown how to combine the algorithms from [13] and [32] to achieve this bound without affecting the other properties shown in [13]; we are confident that the same technique can be applied to the algorithm proposed in this work. Finally, all our algorithms work with any clock drift parameter $1 < \vartheta \leq 1.004$, that is, the nodes' clocks can have up to 0.4% drift. In comparison, cheap quartz oscillators achieve $\vartheta \approx 1 + 10^{-5}$.

1.4 Hardness of Pulse Synchronisation

We consider our results of interest beyond the immediate improvements in complexity of the best known algorithms for pulse synchronisation. Since our framework may employ any consensus algorithm, it proves that pulse synchronisation is, essentially, *as easy* as synchronous consensus—a problem without the requirement for self-stabilisation or any timing uncertainty. Apart from the possibility for future improvements in consensus algorithms carrying over, this accentuates the following fundamental open question:

Is pulse synchronisation *at least as hard* as synchronous consensus?

Due to the various lower bounds and impossibility results on consensus [8, 10, 19, 28] mentioned earlier, a positive answer would immediately imply that the presented techniques are near-optimal. However, one may speculate that pulse synchronisation may rather have the character of (synchronous) approximate agreement [12, 17], as *precise* synchronisation of the pulse events at different nodes is not required. Considering that approximate agreement can be deterministically solved in $O(\log c)$ rounds, where c is the range of the input values, a negative answer is a clear possibility as well. Given that all currently known solutions either explicitly solve consensus, leverage techniques that are likely to be strong enough to solve consensus, or are very slow, this would suggest that new algorithmic techniques and insights into the problem are necessary.

2 PRELIMINARIES

In this section, we describe the model of computation, introduce notation used in the subsequent sections, and formally define the pulse synchronisation and resynchronisation problems.

2.1 Notation

We use $\mathbb{N} = \{1, 2, \dots\}$ to denote positive integers and $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. For any $k \in \mathbb{N}$, we define the shorthand $[k] = \{0, 1, \dots, k-1\}$. Finally, we write $\mathbb{R}^+ = [0, \infty)$ for the set of non-negative real numbers. For $a, b \in \mathbb{R}^+$ we use the notation $[a, b)$ and $(a, b]$ for half-open intervals and $[a, b]$ for closed intervals. Finally, we write $\mathbb{R}^+ \cup \{\infty\} = [0, \infty]$.

2.2 Reference Time and Clocks

Throughout this work, we assume a global *reference time* that is *not* available to the nodes in the distributed system. The reference time is only used to reason about the behaviour of the system. A *clock* is a strictly increasing function $C : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ that maps the reference time to local (perceived) time. That is, at reference time t , clock C indicates that the local time is $C(t)$. We say that a clock C has *drift* at most $\vartheta - 1 > 0$ if for any $t, t' \in \mathbb{R}^+$, where $t < t'$, the clock satisfies

$$t' - t \leq C(t') - C(t) \leq \vartheta(t' - t).$$

That is, if we have two such clocks, then their measurements of elapsed time are at most factor ϑ apart.

2.3 The Bounded-Delay Model

We consider a *bounded-delay message-passing model* of distributed computation. The system is modelled as a fully connected network of n nodes, where V denotes the set of all nodes. We assume that each node has a unique identifier from the set $[n]$. Each node $v \in V$ has local clock $C(v)$ with maximum drift $\vartheta - 1$ for a known global constant $\vartheta > 1$. We assume that the nodes cannot directly read their local clock values, but instead they can set up local timeouts of predetermined length. That is, a node v can request to be signalled after T time units have passed on the node's own local clock $C(v)$ since the timeout was started.

For communication, we assume sender authentication, that is, each node can distinguish the senders of the messages it receives. In other words, every incoming communication link is labelled with the identifier of the sender. Unlike in fully synchronous models, where communication and computation proceeds in lock-step at all nodes, we consider a model in which each message has an associated delay in $(0, d)$. For simplicity, we assume that the *maximum delay* $d \in \Theta(1)$ is a known constant and we consider d as the basic time unit in the system. We note that even though we assume continuous, real-valued clocks, any constant offset in clock readings, e.g., due to discrete clocks, can be modelled by increasing d if needed.

We assume that the system can experience transient faults that arbitrarily corrupt the state of the entire system; we formally define below what this entails in our model. Once the transient faults cease, we assume that up to f of the n nodes in the system may remain Byzantine faulty, that is, they have arbitrary (mis)behaviour and do not necessarily follow the given protocol. We use $F \subseteq V$, where $|F| \leq f$, to denote an arbitrary set of *faulty nodes* and $G = V \setminus F$ is the set of *correct nodes*.

2.4 Algorithms, Configurations, and Executions

Algorithms. We assume that each node executes a finite state machine whose state transitions can depend on the current state of the node, the set of recently received messages, and local timeouts. Formally, an algorithm is a tuple $A = (S, \mathcal{P}, \mathcal{M}, \mathcal{T}, \delta, \mu)$, where

- S is a finite set of *states*,
- $\mathcal{P} \subseteq S$ is a subset of states that trigger a *pulse event*,
- $\mathcal{M}$ is a finite set of *messages*,
- $\mathcal{T} \subseteq \mathbb{R}^+ \times 2^S$ is a finite set of *timers*,
- $\delta : V \times S \times \mathcal{M}^n \times \{0, 1\}^h \rightarrow S$, where $h = |\mathcal{T}|$, is the state transition function, and
- $\mu : V \times V \times S \rightarrow \mathcal{M}$ is a message function.

We now explain in detail how the system state evolves and algorithms operate.

Local Configurations, Timers, and Timeouts. The local configuration $x(v, t)$ of a node $v \in G$ at time $t \in \mathbb{R}^+$ consists of

- (1) its current state $s(v, t) \in S$,
- (2) the state of its *input channels* $m(v, t) \in \mathcal{M}^n$,
- (3) its local clock value $C(v, t) \in \mathbb{R}^+$, and
- (4) timer states $T_k(v, t) \in [0, T_k]$ for each $(T_k, S_k) \in \mathcal{T}$ and $k \in [h]$.

Recall that we assume that the transient faults have left the system in an arbitrary state at time $t = 0$. This entails that for each node $v \in V$ the initial values at $t = 0$ for (1)–(5) are arbitrary.

In the following, we use the shorthand T_k for timer $(T_k, S_k) \in \mathcal{T}$. We say that timer T_k of node v expires at time t if $T_k(v, t)$ changes to 0 at time t . It is expired at time t if $T_k(v, t) = 0$. Timers may cause state transitions of nodes when expiring. Let $e(v, t) \in \{0, 1\}^h$ indicate which timers are

expired, that is, $e_k(v, t) = 1$ if T_k is expired and $e(v, t) = 0$ otherwise. If at time t the value of either $m(v, t)$ changes (that is, the input channels of node $v \in G$ are updated due to a received message) or some local timer expires, the node updates its current state to $s = \delta(v, s', m(v, t), e(v, t))$, where s' is the node's state prior to this computation. If $s \neq s'$, we say that node v transitions to state s at time t (and write $s(v, t) = s$). We remark that this definition allows for the possibility that state transitions happen in arbitrary short succession. However, our algorithms are designed such that only a (small) constant number of transitions is possible in constant time, and computational delays can be treated by interpreting them as part of communication delays.

For convenience, let us define the predicate $\Delta(v, s, t) = 1$ if $v \in G$ transitions to s at time t and $\Delta(v, s, t) = 0$ otherwise. When node v transitions to state $s \in \mathcal{S}$, it resets all timers (T_k, S_k) for which $s \in S_k$. Accordingly, at each time $t > 0$, the timer state is defined as

$$T_k(v, t) = \max\{0, T_k(v, t_{\text{reset}}) - (C(v, t) - C(v, t_{\text{reset}}))\},$$

where t_{reset} is the most recent time node v reset the timer T_k or time 0, that is,

$$t_{\text{reset}} = \max(\{0\} \cup \{t' \leq t : \Delta(v, s, t') = 1, s \in S\}).$$

Note that $T_k(v, t_{\text{reset}}) = T_k$ unless $t_{\text{reset}} = 0$, since with the exception of the arbitrary initial states the timer state is reset to T_k at time t_{reset} .

The bottom line is that, at all times $t \in \mathbb{R}^+$, the timer state $T_k(v, t) \in [0, T_k]$ indicates how much time needs to pass on the local clock $C(v, \cdot)$ of node v until the timer expires; the rather involved definition of how timers behave in order to achieve this property is owed to the requirement of self-stabilisation.

Communication. We say that a node u sends the message $\mu(u, v, s(u, t)) \in \mathcal{M}$ to node v at time t , if the value of $\mu(u, v, s(u, t))$ changes at time t . Moreover, node u is said to *broadcast* the message a at time t if it sends the message a to every v at time t .

As we operate in the bounded-delay setting, sent messages do not arrive at their destinations immediately. To model this, let the *communication delay function* $d_{uv} : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ be a strictly increasing function such that $0 < d_{uv}(t) - t < d$. The input channels of node $u \in G$ satisfy

$$m_v(u, d_{uv}(t)) = \begin{cases} \mu(v, u, s(v, t)) & \text{if } v \in G \\ b(u, v, t) & \text{otherwise,} \end{cases}$$

where $b(u, v, t) \in \mathcal{M}$ is the message a faulty node $v \in F$ decides to transmit to a correct node $u \in G$ at time t . We assume the adversary can freely choose the communication delay functions d_{uv} . Thus, the adversary can control what correct nodes receive from faulty nodes *and* how long the messages sent by correct nodes traverse (up to the maximum delay bound d). Intuitively, $m_v(u, t) \in \mathcal{M}$ denotes the most recent message node v received from node u at time t . Since transient faults may result in arbitrarily corrupted communication channels at time 0, we assume that $m_v(u, t) \in \mathcal{M}$ is arbitrary for $t < d_{uv}(0)$.

The Adversary and Executions. After fixing $f, n \in \mathbb{N}$ and an algorithm A , we assume that an adversary chooses

- (1) the set $F \subseteq V$ of faulty nodes such that $|F| \leq f$,
- (2) the initial local configuration $x(v, 0)$ for all $v \in V$,

and for all $t \in \mathbb{R}^+$ and any $u, v \in V$

- (3) the local clock values $C(v, t)$,
- (4) the message delay functions $d_{uv}(t)$, and
- (5) the messages $b(u, v, t)$ sent by faulty nodes.

Note that if the algorithm A is deterministic, then the adversary's choices for (1)–(5) together with A determine the execution, that is, local configurations $x(v, t)$ for all $v \in G$ and $t \geq 0$. Randomisation may be used in black-box calls to a consensus subroutine only. For brevity, we postpone the discussion of randomisation to Section 7, which covers the results obtained by utilising randomised consensus routines. We remark that minor adjustments to the above definitions may be necessary, depending on the precise model of randomness and power of the adversary; however, this does not affect the reasoning about our framework, which is oblivious to how the employed consensus routine operates.

Logical State Machines and Sliding Window Memory Buffers. For ease of presentation, we do not describe our algorithms in the above low-level state machine formalism, but instead use high-level state machines, where state transitions are conditioned on timer expiration and sliding window memory buffers. While these are not part of the above-described formalism, they are straightforward to implement using additional local timers and states.

Formally, we use a set X of *logical states* and identify each state $s \in S$ with a logical state $\ell(s) \in X$. That is, we have a surjective projection $\ell : S \rightarrow X$ that maps each state s onto its equivalence class $\ell(s)$, i.e., the logical state. In addition, we associate all timers with some logical state, that is, for every $(T_k, S_k) \in \mathcal{T}$, we have that $S_k \in X$ is an equivalence class of states.

We employ *sliding window buffers* in our algorithms. A sliding window buffer of length T stores the set of nodes from which (a certain type of) a message has been received within time T on the node's local clock. Since the local configuration $x(v, 0)$ of a node v is arbitrary at time 0, we have that by time $T + d$ the contents of the sliding window buffer are guaranteed to be valid: if the buffer of $v \in G$ contains a message m from $u \in G$ at time $t \geq T + d$, then u must have sent an m message to v during the interval $(t - T - d, t)$ of reference time. Vice-versa, if u sends a message m at time t to v , the buffer is guaranteed to contain the message during the interval $(t + d, t + T/\vartheta)$ of reference time. We also allow the algorithms to *clear* the sliding window buffers at any point in time by removing all the messages currently contained in the buffer. That is, when a node clears its sliding window buffer at time t , then the buffer contains no message seen before time t .

2.5 Pulse Synchronisation Algorithms

In the pulse synchronisation problem, the task is to have all the correct nodes locally generate pulse events in an almost synchronised fashion, despite arbitrary initial states and the presence of Byzantine faulty nodes. In addition, these pulses have to be well separated. Let $p(v, t) \in \{0, 1\}$ indicate whether a correct node $v \in G$ generates a pulse at time t . Moreover, let $p_k(v, t) \in [t, \infty)$ denote the time when node v generates the k th pulse event at or after time t and $p_k(v, t) = \infty$ if no such time exists. We say that the system has stabilised from time t onwards if

- (1) $p_1(v, t) \leq t + \Phi^+$ for all $v \in G$,
- (2) $|p_k(v, t) - p_k(u, t)| < \sigma$ for all $u, v \in G$ and $k \geq 1$,
- (3) $\Phi^- \leq p_{k+1}(v, t) - \min\{p_k(u, t) : u \in G\} \leq \Phi^+$ for all $v \in G$ and $k \geq 1$,

where Φ^- and Φ^+ are the accuracy bounds controlling the separation of the generated pulses. That is, (1) all correct nodes generate a pulse during the interval $[t, t + \Phi^+]$; (2) the k th pulse of any two correct nodes is less than σ time apart; and (3) for any pair of correct nodes, their subsequent pulses are at least $\Phi^- - \sigma$, but at most Φ^+ time apart.

We say that A is an f -resilient pulse synchronisation algorithm with *skew* σ and *accuracy* $\Phi = (\Phi^-, \Phi^+)$ with stabilisation time $T(A)$, if for any choices of the adversary such that $|F| \leq f$, there exists a time $t \leq T(A)$ such that the system stabilises from time t onwards. This scenario is illustrated in Figure 4.

Fig. 2. An example execution of a resynchronisation algorithm. Eventually, all correct nodes will generate a pulse within a small time window of length ρ . After this, all correct nodes refrain from generating a new pulse for at least Ψ time units.

Finally, we call a pulse synchronisation algorithm $\mathbf{A}$ a T -pulser if the accuracy bounds satisfy $\Phi^-, \Phi^+ \in \Theta(T)$ and $\mathbf{A}$ has skew $\sigma \leq 2d$. We use $M(\mathbf{A})$ to denote the maximum number of bits a correct node sends on a channel per unit time when executing algorithm $\mathbf{A}$.

2.6 Resynchronisation Algorithms

In our pulse synchronisation algorithms, we use so-called resynchronisation pulses to facilitate stabilisation. The resynchronisation pulses are provided by resynchronisation algorithms that solve a weak variant of pulse synchronisation: the guarantee is that eventually all correct nodes generate a single resynchronisation pulse almost synchronously, which is followed by a long period of silence (i.e., no new resynchronisation pulse). At all other times, the behaviour can be arbitrary. See Figure 2 for illustration.

Formally, we say that $\mathbf{B}$ is an f -resilient resynchronisation algorithm with skew ρ and separation window Ψ that stabilises in time $T(\mathbf{B})$, if the following holds: for any choices of the adversary such that $|F| \leq f$, there exists a time $t \leq T(\mathbf{B})$ such that every correct node $v \in G$ locally generates a *resynchronisation pulse* at time $r(v) \in [t, t + \rho)$ and no other resynchronisation pulse before time $t + \rho + \Psi$. We call such a resynchronisation pulse *good*. In particular, we do not impose any restrictions on what the nodes do outside the interval $[t, t + \rho + \Psi)$, that is, there may be *spurious* resynchronisation pulses outside this interval.

Again, we denote by $M(\mathbf{B})$ the maximum number of bits a correct node sends on a channel per unit time when executing $\mathbf{B}$.

2.7 Synchronous Consensus Routines

As we rely on synchronous consensus algorithms, we briefly define the synchronous model of computation for the sake of completeness. In the synchronous model, the computation proceeds in discrete rounds, that is, the nodes have access to a common global clock. In each round $r \in \mathbb{N}$ the nodes (1) send messages based on their current state, (2) receive messages, and (3) perform local computations and update their state for the next round.

In synchronous binary consensus, we assume that each $v \in V$ is given a private input bit $x(v) \in \{0, 1\}$, starts from a fixed initial state (no self-stabilisation), and is to compute output $y(v) \in \{0, 1\}$.

However, there are f Byzantine faulty nodes. An f -resilient synchronous consensus routine C with round complexity $T(C)$ guarantees:

- (1) *Agreement.* There exists $y \in \{0, 1\}$ such that $y(v) = y$ for all $v \in G$.
- (2) *Validity.* If, for $x \in \{0, 1\}$, it holds that $x(v) = x$ for all $v \in G$, then $y = x$.
- (3) *Termination.* Each $v \in G$ decides on $y(v)$ and terminates by round $T(C)$.

We use $M(C)$ to denote the maximum number of bits any $v \in G$ sends to any other node in a single round of any execution of C .

3 THE TRANSFORMATION FRAMEWORK

Our main contribution is a modular framework that allows us to turn any *non-self-stabilising* synchronous consensus algorithm into a self-stabilising pulse synchronisation algorithm in the bounded-delay model. In particular, the transformation yields only a small overhead in time and communication complexity. As our construction is relatively involved, we opt to present it in a top-down fashion. First, we give our main theorem together with its corollaries. Then we state the auxiliary results we need to prove the main theorem and later discuss how these auxiliary results are established.

3.1 The Main Result

Before we formally state our main result, we make the following definition.

Definition 1 (Family of Consensus Routines). Let $R, M, N : \mathbb{N}_0 \rightarrow \mathbb{N}_0$ be functions that satisfy the following conditions:

- (i) for any $f_0, f_1 \in \mathbb{N}$, we have $N(f_0 + f_1) \leq N(f_0) + N(f_1)$, and
- (ii) both $M(f)$ and $R(f)$ are increasing.

We say that $\langle C, R, M, N \rangle$ is a *family of synchronous consensus routines* if for any $f \geq 0$ and $n \geq N(f)$, there exists a synchronous consensus routine $C \in \mathcal{C}$ such that

- C runs on n nodes and is f -resilient,
- each correct node terminates in $T(C) = R(f)$ rounds,
- and each correct node sends at most $M(C) = M(f)$ bits to any other node per round.

Our main technical result states that given such a family of consensus routines, we can obtain pulse synchronisation algorithms with only small additional overhead. We emphasise that the algorithms in $\mathcal{C}$ are not assumed to be self-stabilising.

THEOREM 1. *Let $\langle C, R, M, N \rangle$ be a family of synchronous consensus routines and $f \geq 0$, $n \geq N(f)$, and $1 < \vartheta \leq 1.004$. Then there exists an f -resilient $R(f)$ -pulser A whose stabilisation time $T(A)$ and number of bits $M(A)$ sent over each channel per time unit satisfy*

$$T(A) \in O\left(d + \sum_{k=0}^{\lceil \log f \rceil} R(2^k)\right) \quad \text{and} \quad M(A) \in O\left(1 + \sum_{k=0}^{\lceil \log f \rceil} M(2^k)\right),$$

where the sums are empty when $f = 0$.

In the deterministic case, the *phase king algorithm* [5] provides a family of synchronous consensus routines that satisfy the requirements. Since in the phase king algorithm, all nodes communicate by broadcasts (i.e., send the same information to all other nodes) and the additional communication by our framework satisfies this property as well, the same is true for the derived pulser. Moreover, the phase king protocol achieves optimal resilience [28] with $N(f) = 3f + 1$, constant

message size $M(f) \in O(1)$, and asymptotically optimal [19] round complexity $R(f) \in \Theta(f)$. Thus, this immediately yields the following result.

COROLLARY 1. *For any $f \geq 0$ and $n > 3f$, there is a deterministic f -resilient $\Theta(f)$ -pulser over n nodes that stabilises in $O(f)$ time and has correct nodes broadcast $O(\log f)$ bits per time unit.*

Employing randomised consensus algorithms in our framework is straightforward. We now summarise the main results related to randomised pulse synchronisation algorithms; the details are discussed later in Section 7. First, by applying our construction to a fast and communication-efficient randomised consensus algorithm, e.g., the one by King and Saia [21], we get an efficient randomised pulse synchronisation algorithm.

COROLLARY 2. *Suppose we have private channels. For any $f \geq 0$, constant $\varepsilon > 0$, and $n > (3 + \varepsilon)f$, there is a randomised f -resilient $(\text{polylog } f)$ -pulser over n nodes that stabilises in $\text{polylog } f$ time w.h.p. and has nodes broadcast $\text{polylog } f$ bits per time unit.*

We can also utilise the constant expected time protocol by Feldman and Micali [18]. With some care, we can show that for $R(f) \in O(1)$, Chernoff's bound readily implies that the stabilisation time is not only in $O(\log f)$ in expectation, but also with high probability.

COROLLARY 3. *Suppose we have private channels. For any $f \geq 0$ and $n > 3f$, there is a randomised f -resilient $\Theta(\log f)$ -pulser over n nodes that stabilises in $O(\log f)$ time w.h.p. and has nodes broadcast $\text{poly } f$ bits per time unit.*

3.2 Proof of Theorem 1

The proof of the main result takes an inductive approach. In the inductive step, we assume two pulse synchronisation algorithms with small resilience. We then use these to construct (via some hoops we discuss later) a new pulse synchronisation algorithm with higher resilience. This step is formalised in the following technical lemma, which we prove later.

LEMMA 1. *Let $f, n_0, n_1 \in \mathbb{N}$ and define the values*

$$n = n_0 + n_1, \quad f_0 = \lfloor (f - 1)/2 \rfloor, \quad f_1 = \lceil (f - 1)/2 \rceil.$$

Suppose there exists

- *for both $i \in \{0, 1\}$ an f_i -resilient R -pulser $\mathbf{A}_i$ that runs on n_i nodes with accuracy $\Phi_i = (\Phi_i^-, \Phi_i^+)$ satisfying $\Phi_i^+ / \Phi_i^- \leq \varphi$ for a sufficiently small constant $\varphi > \vartheta$, and*
- *an f -resilient consensus routine $\mathbf{C}$ for a network of n nodes that has running time R and uses messages of at most M bits.*

Then there exists a R -pulser $\mathbf{A}$ that

- *runs on n nodes and has resilience f ,*
- *stabilises in time $T(\mathbf{A}) \in \max\{T(\mathbf{A}_0), T(\mathbf{A}_1)\} + O(R)$,*
- *sends $M(\mathbf{A}) \in \max\{M(\mathbf{A}_0), M(\mathbf{A}_1)\} + O(M)$ bits over each channel per time unit,*
- *has skew $2d$, and*
- *has accuracy bounds Φ^- and Φ^+ that satisfy $\Phi^+ / \Phi^- \leq \varphi$.*

We observe that Theorem 1 is a relatively straightforward consequence of the above lemma.

THEOREM 1. *Let $\langle \mathbf{C}, R, M, N \rangle$ be a family of synchronous consensus routines and $f \geq 0$, $n \geq N(f)$, and $1 < \vartheta \leq 1.004$. Then there exists an f -resilient $R(f)$ -pulser $\mathbf{A}$ whose stabilisation time $T(\mathbf{A})$ and*

number of bits $M(\mathbf{A})$ sent over each channel per time unit satisfy

$$T(\mathbf{A}) \in O\left(d + \sum_{k=0}^{\lceil \log f \rceil} R(2^k)\right) \quad \text{and} \quad M(\mathbf{A}) \in O\left(1 + \sum_{k=0}^{\lceil \log f \rceil} M(2^k)\right),$$

where the sums are empty when $f = 0$.

PROOF. We prove the claim for $f \in \mathbb{N}_0 = \{0\} \cup \bigcup_{k \in \mathbb{N}_0} ([2^k, 2^{k+1}) \cap \mathbb{N})$ using induction on k . As base case, we use $f = 0$. This is trivial for all $n > 0$, as the following algorithm shows. Let $T > 0$ be arbitrary. We can pick a single node as a designated leader who generates a pulse whenever T time units have passed on its local clock. Whenever the leader node pulses, all other nodes observe this within d time units. When the other nodes observe a pulse from the leader, they generate a pulse locally. Thus, for $f = 0$ we obtain a T -pulser that stabilises in $O(d)$ time and sends messages of $O(1)$ bits at most once every $T/\vartheta \in \Theta(T)$ time. Choosing $T = R(0)$ and noting that $R(0) \geq 1$ (even without faults, consensus requires communication if $n > 1$), the claim follows for $f = 0$.

For the inductive step, consider $f \in [2^k, 2^{k+1})$ and suppose that, for all $f' < 2^k$ and $n' \geq N(f')$, there exists an f' -resilient $R(f')$ -pulser algorithm $\mathbf{B}$ on n' nodes with

$$T(\mathbf{B}) \leq \alpha \left(d + \sum_{k'=0}^{\lceil \log f' \rceil} R(2^{k'}) \right) \quad \text{and} \quad M(\mathbf{B}) \leq \beta \left(1 + \sum_{k'=0}^{\lceil \log f' \rceil} M(2^{k'}) \right),$$

where α and β are sufficiently large constants. In particular, we can now apply Lemma 1 with $f_0, f_1 \leq f/2 < 2^k$ and any $n \geq N(f)$, as $N(f) \geq N(f_0) + N(f_1)$ guarantees that we may choose some $n_0 \geq N(f_0)$ and $n_1 \geq N(f_1)$ such that $n = n_0 + n_1$. This yields an f -resilient $R(f)$ -pulser $\mathbf{A}$ over n nodes, with stabilisation time

$$\begin{aligned} T(\mathbf{A}) &\leq \max\{T(\mathbf{A}_0), T(\mathbf{A}_1)\} + \gamma R(f) \\ &\leq \alpha \left(d + \sum_{k=0}^{\lceil \log f/2 \rceil} R(2^k) \right) + \gamma R(2^{\lceil \log f \rceil}), \\ &\leq \alpha \left(d + \sum_{k=0}^{\lceil \log f \rceil} R(2^k) \right), \end{aligned}$$

where $\gamma \leq \alpha$ is a constant; the second step uses that $R(f)$ is increasing. Similarly, the bound on the number of sent bits follows from Lemma 1 and the induction assumption:

$$\begin{aligned} M(\mathbf{A}) &\leq \max\{M(\mathbf{A}_0), M(\mathbf{A}_1)\} + \gamma' M(f) \\ &\leq \beta \left(1 + \sum_{k=0}^{\lceil \log f/2 \rceil} M(2^k) \right) + \gamma' M(2^{\lceil \log f \rceil}), \\ &\leq \beta \left(1 + \sum_{k=0}^{\lceil \log f \rceil} M(2^k) \right), \end{aligned}$$

where $\gamma' \leq \beta$ is a constant and we used that $M(f)$ is increasing. $\square$

3.3 The Auxiliary Results

In order to show Lemma 1, we use two main ingredients: (1) a pulse synchronisation algorithm whose stabilisation mechanism is triggered by a resynchronisation pulse and (2) a

resynchronisation algorithm providing the latter. These ingredients are formalised in the following two theorems which are proven in Sections 5 and 6, respectively.

THEOREM 2. *Let $f \geq 0$, $n > 3f$, and $(2 + \sqrt{32})/7 > \vartheta > 1$. Suppose for a network of n nodes, there exists*

- *an f -resilient synchronous consensus algorithm $\mathbf{C}$, and*
- *an f -resilient resynchronisation algorithm $\mathbf{B}$ with skew $\rho \in O(d)$ and sufficiently large separation window $\Psi \in O(R)$ that tolerates clock drift of ϑ ,*

where $\mathbf{C}$ runs in $R = R(f)$ rounds and lets nodes send at most $M = M(f)$ bits per round and channel. Then there exists $\varphi_0(\vartheta) \in 1 + O(\vartheta - 1)$ such that, for any constant $\varphi > \varphi_0(\vartheta)$ and sufficiently large $T \in O(R)$, there exists an f -resilient pulse synchronisation algorithm $\mathbf{A}$ for n nodes that

- *has skew $\sigma = 2d$,*
- *satisfies the accuracy bounds $\Phi^- = T$ and $\Phi^+ = T\varphi$,*
- *stabilises in $T(\mathbf{B}) + O(R)$ time, and*
- *has nodes send $M(\mathbf{B}) + O(M)$ bits per time unit and channel.*

To apply the above theorem, we require suitable consensus and resynchronisation algorithms. We rely on consensus algorithms from prior work and construct efficient resynchronisation algorithms ourselves. The idea is to combine pulse synchronisation algorithms that have *low resilience* to obtain resynchronisation algorithms with *high resilience*.

THEOREM 3. *Let $f, n_0, n_1 \in \mathbb{N}$ and $1 < \vartheta \leq 1.004$. Define*

$$n = n_0 + n_1, \quad f_0 = \lfloor (f - 1)/2 \rfloor, \quad f_1 = \lceil (f - 1)/2 \rceil.$$

For any $\Psi \in \Omega(1)$ and sufficiently small constant $\varphi > \varphi_0(\vartheta)$, there exists a bound $T_0 \in \Theta(\Psi)$ such that the following claim holds. If, for both $i \in \{0, 1\}$, there exists pulse synchronisation algorithm $\mathbf{A}_i$ that

- *runs on n_i nodes and has resilience f_i ,*
- *has skew $\sigma = 2d$, and*
- *has accuracy bounds $\Phi_i^- = T$ and $\Phi_i^+ = T\varphi$, where $T_0 \leq T$ and $T \in O(\Psi)$,*

then there exists a resynchronisation algorithm $\mathbf{B}$ that

- *runs on n nodes and has resilience f ,*
- *has skew $\rho \in O(d)$ and separation window of length Ψ ,*
- *generates a resynchronisation pulse by time $T(\mathbf{B}) \in \max\{T(\mathbf{A}_0), T(\mathbf{A}_1)\} + O(\Psi)$, and*
- *has nodes send $M(\mathbf{B}) \in \max\{M(\mathbf{A}_0), M(\mathbf{A}_1)\} + O(1)$ bits per time unit and channel.*

Given a suitable consensus algorithm, one can readily combine Theorems 2 and 3 to obtain Lemma 1. Note that for both theorems, it turns out that $\varphi_0(\vartheta) = 1 + 5(\vartheta - 1)/(2 + 2\vartheta - 3\vartheta^2)$ will do. Therefore, we can reduce the problem of constructing an f -resilient pulse synchronisation algorithm to finding algorithms that tolerate up to $\lfloor f/2 \rfloor$ faults and recurse; Figure 3 illustrates how these two types of algorithms are interleaved.

3.4 Proof of Lemma 1

PROOF. From Theorem 3, we get that for any sufficiently large $\Psi \in \Theta(R)$, there exists a resynchronisation algorithm $\mathbf{B}$ with skew $\rho \in O(d)$ and separation window of length Ψ that

- runs on n nodes and has resilience f ,
- stabilises in time $\max\{T(\mathbf{A}_0), T(\mathbf{A}_1)\} + O(\Psi) = \max\{T(\mathbf{A}_0), T(\mathbf{A}_1)\} + O(R)$, and
- has nodes send $\max\{M(\mathbf{A}_0), M(\mathbf{A}_1)\} + O(1)$ bits per time unit and channel.

Fig. 3. Recursively building a 2-resilient pulse synchronisation algorithm $A(7, 2)$ over 7 nodes. The construction utilises low resilience pulse synchronisation algorithms to build high resilience resynchronisation algorithms which can then be used to obtain highly resilient pulse synchronisation algorithms. Here, the base case consists of trivial 0-resilient pulse synchronisation algorithms $A(2, 0)$ and $A(3, 0)$ over two and three nodes, respectively. Two copies of $A(2, 0)$ are used to build a 1-resilient resynchronisation algorithm $B(4, 1)$ over 4 nodes using Theorem 3. The resynchronisation algorithm $B(4, 1)$ is used together with a synchronous consensus algorithm $C(4, 1)$ to obtain a pulse synchronisation algorithm $A(4, 1)$ via Theorem 2. Now, the 1-resilient pulse synchronisation algorithm $A(4, 1)$ over 4 nodes is used together with the trivial 0-resilient algorithm $A(3, 0)$ to obtain a two-resilient resynchronisation algorithm $B(7, 2)$ for 7 nodes. This is then used together with a 2-resilient consensus algorithm $C(7, 2)$ to obtain the final pulse synchronisation algorithm $A(7, 2)$. White nodes represent correct nodes and black nodes represent faulty nodes. The gray blocks contain too many faulty nodes for the respective algorithms to correctly operate, and hence, they may have arbitrary output.

We feed B and C into Theorem 2, yielding a pulse synchronisation algorithm A with the claimed properties, as the application of Theorem 2 increases the stabilisation time by an additional $O(R)$ time units and adds $O(M)$ bits per time unit and channel. $\square$

3.5 Organisation of the Remainder of the Article

We dedicate the remaining sections to fill in the details we have omitted above. Namely,

- Section 4 describes a Byzantine-tolerant pulse synchronisation algorithm that is *not* self-stabilising. We utilise the algorithm in Section 5, but the section also serves to provide a gentle introduction to the notation and style of proofs we use in the following sections;
- Section 5 gives the proof of Theorem 2;
- Section 6 gives the proof of Theorem 3; and
- Section 7 extends our framework to operate with randomised consensus algorithms. This establishes Corollaries 2 and 3.

4 BYZANTINE-TOLERANT PULSE SYNCHRONISATION

In this section, we describe a *non-self-stabilising* pulse synchronisation algorithm, which we utilise later in our construction of the self-stabilising algorithm. The algorithm given here is a variant of the Byzantine fault-tolerant clock synchronisation algorithm by Srikanth and Toeug [31] that avoids transmitting clock values in favour of unlabelled pulses.

Fig. 4. An example execution of a non-self-stabilising pulse synchronisation algorithm with initialisation signals.

4.1 Pulse Synchronisation without Self-Stabilisation

As previously mentioned, we do not require self-stabilisation for now. However, instead of assuming a synchronous start, we devise an algorithm where all correct nodes synchronously start generating pulses once all correct nodes have received an *initialisation signal* within a short time window. We will later show how to generate such initialisation signals in a self-stabilising manner. In particular, we allow that *before* receiving an initialisation signal, a correct node can have arbitrary behaviour, but *after* a correct node has received this signal, it waits until all correct nodes have received the signal and start to synchronously generate well-separated pulses, as shown in Figure 4.

In the following, suppose that all correct nodes receive an initialisation signal during the time window $[0, \tau)$. In other words, nodes can start executing the algorithm at different times, but they all do so by some bounded (possibly non-constant) time τ . When a node receives the initialisation signal, it immediately transitions to a special RESET state, whose purpose is to consistently initialise local memory and wait for other nodes to receive the initialisation signal as well; before this, correct nodes can have arbitrary behaviour.

Later, we will repeatedly make use of this algorithm as a subroutine for a self-stabilising algorithm and we need to consider the possibility that there are still messages from earlier (possibly corrupted) instances in transit, or nodes may be executing a previous instance in an incorrect way. Given the initialisation signal, this is easily overcome by waiting for sufficient time before leaving the starting state: waiting $\vartheta(\tau + d) \in O(\tau)$ local time guarantees that (i) all correct nodes transitioned to the starting state and (ii) all messages sent before these transitions have arrived. Clearing memory buffers when leaving the starting state thus ensures that no obsolete information from previous instances is stored by correct nodes.

The goal of this section is to establish the following theorem:

THEOREM 4. *Let $n > 1$, $f < n/3$, and $\tau > 0$. If every correct node receives an initialisation signal during $[0, \tau)$, then there exists a pulse synchronisation algorithm $\mathbf{P}$ that*

- runs on n nodes and has resilience f ,
- has $v \in G$ generate its first pulse (after the initialisation signal) at a time $t_0(v) \in O(\vartheta^2 d \tau)$,
- has skew $2d$,
- has accuracy bounds $\Phi^- \in \Omega(\vartheta d)$ and $\Phi^+ \in O(\vartheta^2 d)$, and
- lets each node broadcast at most one bit per time unit.

4.2 Description of the Algorithm

The algorithm is illustrated in Figure 5. In the figure, the circles denote the basic logical states (RESET, START, READY, PROPOSE, PULSE) of the state machine for each node. The two states RESET and START are used in the initialisation phase of the algorithm, which takes place during $[0, \tau)$ when

Fig. 5. The state machine for the non-self-stabilising pulse synchronisation algorithm. State transitions occur when the condition of the guard in the respective edge is satisfied (labelled gray boxes). Here, all transition guards involve checking whether a local timer expires or a node has received sufficiently many messages from nodes in state PROPOSE. The only communication that occurs is when a node transitions to state PROPOSE; when this happens a node broadcasts this information to all others. The notation $\langle T \rangle$ indicates the expiration of a timer of length T that was reset when transitioning to the current state, that is, T time units have passed on the local clock since the transition to the current state. The box labelled PROPOSE indicates that a node clears its sliding window messages buffers when transitioning from RESET to START and PULSE to READY. That is, the node forgets who it has “seen” in PROPOSE, the previous iteration. The algorithm assumes that during the interval $[0, \tau]$ all nodes transition to RESET. This starts the initialisation phase of the algorithm. Eventually, all nodes transition to PULSE within a short time window and start executing the algorithm. Whenever a node transitions to state PULSE, it generates a local pulse event. Table 2 lists the constraints imposed on the timeouts.

the nodes receive their initialisation signals. In Figure 5, directed edges between the states denote possible state transitions and labels give the conditions (transition guards) when the transition is allowed to (and must) occur. The notation is as follows: $\langle T_k \rangle$ denotes the condition that timer $(T_k, \{s\})$ has expired, where s is the state that is left when the guard is satisfied.

The boxes labelled with PROPOSE indicate that when a node transitions to the designated state, it clears its memory buffers immediately. Throughout this section, we use $H(v, t)$ to denote the total number of nodes at time t from which v has received a PROPOSE message since last clearing its memory buffers. For the purposes of our analysis, we use $K(v, t)$ to denote the number of *correct* nodes from which v has received a PROPOSE message at time t since last clearing its sliding window buffers. Moreover, without loss of generality, we may assume that the sliding window buffers have infinite length, that is, messages never expire unless the buffer is explicitly cleared during a transition to START and READY.

The constraints we impose on the timeouts are given in Table 2. The expression “ $> f$ propose messages received” denotes the condition $H(v, t) > f$. For simplicity, we assume in all our descriptions that every time a node transitions to a logical state, it broadcasts the name of the state to all other nodes. Given that we use a constant number of (logical) states per node (and in our algorithms, nodes can only undergo a constant number of state transitions in constant time), this requires broadcasting $O(1)$ bits of information per time unit. In fact, closer inspection reveals that 1 bit per iteration of the cycle suffices here: the only relevant information is whether a node is in state PROPOSE or not.

4.3 Algorithm Analysis

The algorithm relies heavily on the property that there are at most $f < n/3$ faulty nodes. This allows the use of the following “vote-and-pull” technique. If some correct node receives a PROPOSE

Table 2. The List of Conditions
Used in the Non-self-stabilising
Pulse Synchronisation
Algorithm Given in Figure 5

(1)	$T_0/\vartheta \geq \tau + d$
(2)	$T_1/\vartheta \geq (1 - 1/\vartheta)T_0 + \tau$
(3)	$T_2/\vartheta \geq 3d$
(4)	$T_3/\vartheta \geq (1 - 1/\vartheta)T_2 + 2d$

Recall that $d, \vartheta \in O(1)$ and τ is a parameter of the algorithm.

message from at least $n - f$ different nodes at time t , then we must have that at least $n - 2f > f$ of these originated from correct nodes during the interval $(t - d, t)$, as every message has a positive delay of less than d . Furthermore, it follows that before time $t + d$ all correct nodes receive more than f PROPOSE messages.

In particular, this “vote-and-pull” technique is used in the transition to states PROPOSE and PULSE. Suppose at some point all nodes are in READY. If some node transitions to PULSE, then it must have observed at least $n - f$ nodes in PROPOSE by Guard G3. This in turn implies that more than f correct nodes have transitioned to PROPOSE. This in turn will (in short time) “pull” nodes that still remain in state READY into state PROPOSE. Thus, Guard G3 will eventually be satisfied at all the nodes. The same technique is also used in the transition from START to PROPOSE during the initialisation phase of the algorithm.

Remark 1. Suppose $f < n/3$. Let $u, v \in G$, $t \geq d$ and $I = (t - d, t + d)$. If $H(v, t) \geq n - f$, then $K(u, t') > f$ for some $t' \in I$ assuming u does not clear its message buffers during the interval I .

For all $v \in G$ and $t > 0$, let $p(v, t) \in \{0, 1\}$ indicate whether v transitions to state PULSE at time t . That is, we have $p(v, t) = 1$ if node $v \in G$ transitions to state PULSE at time t and $p(v, t) = 0$ otherwise.

LEMMA 2. *There exists $t_0 < \tau + T_0 + T_1 + d$ such that for all $v \in G$ it holds that $p(v, t) = 1$ for $t \in [t_0, t_0 + 2d)$.*

PROOF. Let $v \in G$. Node v receives the initialisation signal during some time $t_{\text{reset}}(v) \in [0, \tau)$ and transitions to state RESET. From RESET, the node transitions to START at some time $t_{\text{start}}(v) \in [t_{\text{reset}}(v) + T_0/\vartheta, t_{\text{reset}}(v) + T_0]$ when the timer T_0 in Guard G1 expires. Since $T_0/\vartheta \geq \tau + d$ by Constraint (1), we get that $t_{\text{start}}(v) \geq \tau + d$. Thus, for all $u, v \in G$, we have $t_{\text{reset}}(u) + d \leq t_{\text{start}}(v)$.

Moreover, v transitions to PROPOSE at some time $t_{\text{propose}}(v) \in [t_{\text{start}}(v), t_{\text{start}}(v) + T_1]$ when Guard G2 is satisfied. Hence, any $v \in G$ transitions to state PROPOSE no later than time $t_{\text{start}}(v) + T_1 \leq t_{\text{reset}}(v) + T_0 + T_1 \leq \tau + T_0 + T_1$. Let $t_{\text{propose}} \geq \tau + d$ be the minimal time some node $v \in G$ transitions to state PROPOSE after transitioning to RESET during $[0, \tau)$. Observe that since a correct node v clears its message buffers when transitioning from RESET to START, we have that, for any $t \in [t_{\text{start}}(v), t_{\text{propose}}] \subseteq [\tau + d, t_{\text{propose}})$, the sliding window memory buffer of v contains no messages from correct nodes at time t , i.e., $K(v, t) = 0$ and $H(v, t) \leq f$. Thus, node v will not receive a PROPOSE message from any correct node $u \in G$ before time t_{propose} .

Note that $t_{\text{propose}} \in [(T_0 + T_1)/\vartheta, \tau + T_0 + T_1]$ by Guard G1 and Guard G2. By Constraint (2) and our previous bounds, we have that $t_{\text{propose}} \geq T_0/\vartheta + (1 - 1/\vartheta)T_0 + \tau = T_0 + \tau \geq t_{\text{start}}(u)$ for any $u \in G$. Hence, after time $T_0 + \tau$, no $u \in G$ clears its memory buffer before transitioning to PULSE at time $t_{\text{pulse}}(u)$. In particular, we now have that $t_{\text{propose}}(v) \leq \tau + T_0 + T_1$ and hence all nodes transition

to PULSE by some time $t_{\text{pulse}}(v) < \tau + T_0 + T_1 + d$, as each $u \in G$ must have received a PROPOSE message by this time from least $n - f$ correct nodes, meeting the condition of Guard G3.

Let $t_0 = \min\{t_{\text{pulse}}(v) : v \in G\} < \tau + T_0 + T_1 + d$ be the minimal time some correct node transitions to state PULSE. It remains to argue that $t_{\text{pulse}}(v) \in [t_0, t_0 + 2d]$. By Constraint (3), no correct node clears its memory buffer before time $t_0 + 3d$. Since some node $v \in G$ transitioned to PULSE at time t_0 , we must have that its condition in Guard G3 was satisfied. That is, node v must have received a PROPOSE message from at least $n - f$ nodes since clearing its memory buffer at time $t_{\text{start}}(v)$, that is, $H(v, t_0) \geq n - f$ and thus $K(v, t_0) > f$.

As these messages must have been received after time t_{propose} , by then each $u \in G$ already reached state START and, by Constraint (3), no correct node can reset its PROPOSE flags again before time $t_0 + 3d$, it follows that $K(u, t_0 + d) > f$ for each $u \in G$. In particular, each $u \in G$ transitions to state PROPOSE by time $t_0 + d$. It now follows that, at time $t'_0 < t_0 + 2d$, we have $K(u, t'_0) \geq n - f$ for all $u \in G$, implying that Guard G2 is satisfied for each such u . Thus, $t_{\text{pulse}}(u) \in [t_0, t'_0] \subseteq [t_0, t_0 + 2d]$ for each $u \in G$, as claimed. $\square$

Let us now fix t_0 as given by the previous lemma. For every correct node $v \in G$, we define

$$p_0(v) = \inf\{t \geq t_0 : p(v, t) = 1\} \quad \text{and} \quad p_{i+1}(v) = p_{\text{next}}(v, p_i(v)),$$

where $p_{\text{next}}(v, t) = \inf\{t' > t : p(v, t') = 1\}$ is the next time after time t node v generates a pulse.

LEMMA 3. *For all $i \geq 0$, there exist*

$$t_{i+1} \in [t_i + (T_2 + T_3)/\vartheta, t_i + T_2 + T_3 + 3d] \quad \text{such that} \quad p_i(v) \in [t_i, t_i + 2d] \text{ for all } v \in G.$$

PROOF. We show the claim using induction on i . For the case $i = 0$, the claim $p_0(v) \in [t_0, t_0 + 2d]$ follows directly from Lemma 2. For the inductive step, suppose $p_i(v) \in [t_i, t_i + 2d]$ for all $v \in G$. Each $v \in G$ transitions to state READY at a time $t_{\text{ready}}(v) \in [t_i + T_2/\vartheta, t_i + 2d + T_2]$ by Guard G4. Moreover, by Constraint (4), we have that $t_{\text{ready}}(v) > t_i + T_2/\vartheta \geq t_i + 3d$. As no correct node transitions to PROPOSE during $[t_i + 2d, t_i + (T_2 + T_3)/\vartheta]$, this implies that no node receives a PROPOSE message from a correct node before the time $t_{\text{propose}}(u)$ when some node u transitions to PROPOSE from READY (for the next time after $t_i + 3d$). Observe that $t_{\text{propose}}(u) > t_i + (T_2 + T_3)/\vartheta > t_i + 2d + T_2$ by Guard G5 and Constraint (4). Thus, we have $t_{\text{ready}}(v) < t_i + 2d + T_2 < t_{\text{propose}}(u)$ for all $u, v \in G$. Therefore, there exists a time $t_{\text{ready}} < t_i + 2d + T_2$ such that all correct nodes are in state READY and $K(v, t_{\text{ready}}) = 0$ for all $v \in G$.

Next observe that $t_{\text{propose}}(v) \leq t_i + 2d + T_2 + T_3$ for any $v \in G$. Hence, every $u \in G$ will receive a PROPOSE message from every $v \in G$ before time $t_{\text{propose}}(v) + d \leq t_i + 3d + T_2 + T_3$. Thus, by Guard G3 we have that u transitions to PULSE yielding that $p_{i+1}(v) \in [t_i + t_{\text{ready}}, t_i + 3d + T_2 + T_3] \subseteq [t_i + (T_2 + T_3)/\vartheta, t_i + T_2 + T_3 + 3d]$. Let $t_{i+1} = \inf\{p_{i+1}(v) : v \in G\}$. We have already established that $t_{i+1} \in [t_i + (T_2 + T_3)/\vartheta, t_i + T_2 + T_3 + 3d]$. Now using the same arguments as in Lemma 2, it follows that for each $u \in G$, $t_{\text{propose}}(u) < t_{i+1} + d$, as u must have received more than f PROPOSE messages before time $t_{i+1} + d$ triggering the condition in Guard G5 for node u . Thus, Guard G3 will be satisfied before time $t_{i+1} + 2d$ at each $u \in G$, implying that $p_{i+1}(u) \in [t_{i+1}, t_{i+1} + 2d]$ for each $u \in G$. $\square$

THEOREM 4. *Let $n > 1$, $f < n/3$, and $\tau > 0$. If every correct node receives an initialisation signal during $[0, \tau)$, then there exists a pulse synchronisation algorithm **P** that*

- runs on n nodes and has resilience f ,
- has $v \in G$ generate its first pulse (after the initialisation signal) at a time $t_0(v) \in O(\vartheta^2 d \tau)$,
- has skew $2d$,
- has accuracy bounds $\Phi^- \in \Omega(\vartheta d)$ and $\Phi^+ \in O(\vartheta^2 d)$, and
- lets each node broadcast at most one bit per time unit.

PROOF. The constraints in Table 2 are satisfied by setting

$$\begin{aligned} T_0 &= \vartheta(\tau + d) \\ T_1 &= \vartheta^2(1 - 1/\vartheta)(\tau + d) + \tau \\ T_2 &= \vartheta 3d \\ T_3 &= \vartheta^2(1 - 1/\vartheta)3d + 2d. \end{aligned}$$

By Lemma 2, we get that there exists $t_0 \in O(\vartheta^2 d \tau)$ such that all nodes generate the first pulse during the interval $[t_0, t_0 + 2d)$. Applying Lemma 3, we get that, for all $i > 0$, we have that nodes generate the i th pulse during the interval $[t_i, t_i + 2d)$, where $t_i \in [t_{i-1} + (T_2 + T_3)/\vartheta, t_{i-1} + T_2 + T_3 + 3d) \subseteq [\Phi^-, \Phi^+)$. Note that $T_2 + T_3 \in \Theta(\vartheta^2 d)$ and $\vartheta, d \in O(1)$. These observations give the first four properties. For the final property, observe that nodes only need to communicate when they transition to PROPOSE. By Guard G4, correct nodes wait at least for $T_2/\vartheta = 3d$ reference time before transitioning to PROPOSE again after generating a pulse. Hence, nodes need to broadcast at most one bit every $3d > d$ time. $\square$

5 SELF-STABILISING PULSE SYNCHRONISATION

In this section, we show how to use a resynchronisation algorithm and a synchronous consensus routine to devise *self-stabilising* pulse synchronisation algorithms. We obtain the following result:

THEOREM 2. *Let $f \geq 0$, $n > 3f$, and $(2 + \sqrt{32})/7 > \vartheta > 1$. Suppose for a network of n nodes, there exists*

- *an f -resilient synchronous consensus algorithm C , and*
- *an f -resilient resynchronisation algorithm B with skew $\rho \in O(d)$ and sufficiently large separation window $\Psi \in O(R)$ that tolerates clock drift of ϑ ,*

where C runs in $R = R(f)$ rounds and lets nodes send at most $M = M(f)$ bits per round and channel. Then there exists $\varphi_0(\vartheta) \in 1 + O(\vartheta - 1)$ such that, for any constant $\varphi > \varphi_0(\vartheta)$ and sufficiently large $T \in O(R)$, there exists an f -resilient pulse synchronisation algorithm A for n nodes that

- *has skew $\sigma = 2d$,*
- *satisfies the accuracy bounds $\Phi^- = T$ and $\Phi^+ = T\varphi$,*
- *stabilises in $T(B) + O(R)$ time, and*
- *has nodes send $M(B) + O(M)$ bits per time unit and channel.*

5.1 Overview of the Ingredients

The pulse synchronisation algorithm presented in this section consists of two state machines running in parallel:

- (1) the so-called *main state machine* that is responsible for pulse generation, and
- (2) an *auxiliary state machine*, which assists in initiating consensus instances and stabilisation.

The main state machine indicates when pulses are generated and handles all the communication between nodes except for messages sent by simulated consensus instances. The latter are handled by the auxiliary state machine. The transitions in the main state machine are governed by a series of threshold votes, local timeouts, and signals from the auxiliary state machine.

As we aim to devise self-stabilising algorithms, the main and auxiliary state machines may be arbitrarily initialised. To handle this, a stabilisation mechanism is used in conjunction to ensure that, regardless of the initial state of the system, all nodes eventually manage to synchronise their

Fig. 6. Constructing a self-stabilising (SS) and Byzantine fault-tolerant pulse synchronisation algorithm A in the bounded-delay model (BD) out of Byzantine fault-tolerant non-stabilising pulse synchronisation algorithm P, synchronous consensus algorithm C, and resynchronisation algorithm B. All algorithms run on the same set of nodes. (1) The resynchronisation algorithm B eventually outputs a good resynchronisation pulse, which resets the stabilisation mechanism used by the auxiliary state machine. (2) The auxiliary state machine simulates the executions of C using P. Simulations are initiated either due to nodes transitioning to a special WAIT state of the main state machine (see Figure 7) or a certain time after a resynchronisation pulse. (3) The main state machine. It generates pulses when a consensus instance outputs “1” and, when stabilised, guarantees re-initialisation of the consensus algorithm by the auxiliary state machine.

state machines. The stabilisation mechanism relies on the following three subroutines which are summarised in Figure 6:

- (a) a resynchronisation algorithm B,
- (b) the non-self-stabilising pulse synchronisation algorithm P from Section 4, and
- (c) a synchronous consensus algorithm C.

Resynchronisation Pulses. Recall that *resynchronisation algorithm B* solves a weak variant of a pulse synchronisation: it guarantees that *eventually*, within some bounded time $T(B)$, all correct nodes generate a *good* resynchronisation pulse such that no new resynchronisation pulse is generated before Ψ time has passed. Note that, at all other times, algorithm B is allowed to generate pulses at arbitrary frequencies and not necessarily at all correct nodes. Nevertheless, at some point, all correct nodes are bound to generate a good resynchronisation pulse in rough synchrony. We leverage this property to cleanly re-initialise the stabilisation mechanism from time to time. Observe that the idea is somewhat similar to the use of initialisation signals in Section 4, but now we have less control over the incoming “initialisation signals” (i.e., resynchronisation pulses). To summarise, we assume throughout this section that every correct node $v \in G$

- receives a single resynchronisation pulse at time $t_v \in [0, \rho)$, and
- does not receive another resynchronisation pulse before time $t_v + \Psi$,

where Ψ is sufficiently large value we determine later. Later in Section 6, we devise efficient algorithms that produce the needed resynchronisation pulses.

Simulating Synchronous Consensus. The two subroutines (b) and (c) are used in conjunction as follows. We use the variant of the Srikanth-Toueg pulse synchronisation algorithm P described in Section 4 to simulate a synchronous consensus algorithm (in the bounded-delay model). Note that while this pulse synchronisation algorithm is not self-stabilising, it works properly even if non-faulty nodes initialise the algorithm at different times as long as they do so within a time interval of length τ .

Assuming that the nodes initialise the non-self-stabilising pulse synchronisation algorithm P within at most time τ apart, it is straightforward to simulate round-based (i.e., synchronous)

Guard	Condition
G1	$\langle T_1 \rangle$ expires and received $\geq n - f$ PULSE messages within time T_1 before $\langle T_1 \rangle$ expired
G1'	$\langle T_1 \rangle$ expires and $\neg G1$
G2	auxiliary machine signals OUTPUT 1
G2'	$\langle T_{\text{wait}} \rangle$ expires or auxiliary machine signals OUTPUT 0

Fig. 7. The main state machine. When a node transitions to state PULSE (double circle) it will generate a local pulse event and send a PULSE message to all nodes. When the node transitions to state WAIT it broadcasts a WAIT message to all nodes. Guard G1 employs a sliding window memory buffer, which stores any PULSE messages that have arrived within time T_1 (as measured by the local clock). When a correct node transitions to PULSE it resets a local T_1 timeout. Once this expires, either Guard G1 or Guard G1' become satisfied. Similarly, the timer T_{wait} is reset when the node transitions to WAIT. Once it expires, Guard G2' is satisfied and the node transitions from WAIT to RECOVER. The node can transition to PULSE state when Guard G2 is satisfied, which requires an OUTPUT 1 signal from the auxiliary state machine given in Figure 8.

algorithms: a pulse generated by P indicates that a new round of the synchronous algorithm can be started. By setting the delay between two pulses large enough, we can ensure that

- (1) all nodes have time to execute the local computations of the synchronous algorithm, and
- (2) all messages related to a single round arrive before a new pulse occurs.

Employing Silent Consensus. We utilise so-called silent consensus routines in our construction. Silent consensus routines satisfy exactly the same properties as usual consensus routines (validity, agreement, and termination) with the addition that correct nodes send no messages in executions in which all nodes have input 0.

Definition 2 (Silent consensus). A consensus routine is *silent*, if in each execution in which all correct nodes have input 0, correct nodes send no messages.

Any synchronous consensus routine can be converted into a silent consensus routine essentially for free. In our prior work [24], we showed that there exists a simple transformation that induces only an overhead of two rounds while keeping all other properties of the algorithm the same.

THEOREM 5 ([24]). *Any consensus protocol C that runs in R rounds can be transformed into a silent consensus protocol C' that runs in R + 2 rounds. Moreover, the resilience and message size of C and C' are the same.*

Thus, without loss of generality, we assume throughout this section that the given consensus routine C is silent. Moreover, this does not introduce any asymptotic loss in the running time or number of bits communicated.

5.2 High-level Idea of the Construction

The high-level strategy used in our construction is as follows: We run the resynchronisation algorithm in parallel to the self-stabilising pulse synchronisation algorithm we devise in this section. The resynchronisation algorithm will send the resynchronisation signals it generates to the pulse synchronisation algorithm as shown in Figure 6.

The pulse synchronisation algorithm consists of the main state machine given in Figure 7 and the auxiliary state machine given in Figure 8. The auxiliary state machine is responsible for generating the output signals that drive the main state machine (Guard G2 and Guard G2').

Fig. 8. The auxiliary state machine. The auxiliary state machine is responsible for initialising and simulating the consensus routine. The gray boxes denote states which represent the simulation of the consensus routine C. When transitioning to either RUN 0 or RUN 1, the node locally initialises the (non-self-stabilising) pulse synchronisation algorithm from Section 4 and a new instance of C. If the node transitions to RUN 0, it uses input 0 for the consensus routine. If the node transitions to RUN 1, it uses input 1. When the consensus simulation declares an output, the node transitions to either OUTPUT 0 or OUTPUT 1 (sending the respective output signal to the main state machine) and immediately to state LISTEN. The timeouts T_{listen} , T_2 , and $T_{\text{consensus}}$ are reset when a node transitions to the respective states that use a guard referring to them. The timeout T_{active} in Guard G3 (dashed line) is reset by the resynchronisation signal from the underlying resynchronisation algorithm B. Both INPUT 0 and INPUT 1 have a self-loop that is activated if Guard G4 is satisfied. This means that if Guard G4 is satisfied while in these states, the timer T_2 is reset.

The auxiliary state machine employs a consensus routine to facilitate agreement among the nodes on whether a new pulse should occur. If the consensus simulation outputs 1 at some node, then the auxiliary state machine signals the main state machine to generate a pulse. Otherwise, if the consensus instance outputs 0, then this is used to signal that something is wrong and the node can detect that the system has not stabilised. We carefully set up our construction so that once the system stabilises, any consensus instance run by the nodes is guaranteed to always output 1 at every correct node.

As we operate under the assumption that the initial state is arbitrary, the non-trivial part in our construction is to get all correct nodes synchronised well enough to even start simulating consensus jointly in the first place. This is where the resynchronisation algorithm comes into play. We make sure that the algorithm either stabilises or all nodes get “stuck” in a recovery state RECOVER. To deal with the latter case, we use the resynchronisation pulse to let all nodes synchronously reset a local timeout. Once this timeout expires, nodes that are in state RECOVER start a consensus instance with input “1”. By the time this happens, either

- the algorithm has already stabilised (and thus no correct node is in state RECOVER), or
- all correct nodes are in state RECOVER and jointly start a consensus instance that will output “1” (by validity of the consensus routine).

In both cases, stabilisation is guaranteed.

Receiving a Resynchronisation Signal. The use of the resynchronisation signal is straightforward: when a correct node $u \in G$ receives a resynchronisation signal from the underlying resynchronisation algorithm B, node u resets its local timeout T_{active} used by the auxiliary state machine in Figure 8. Upon expiration of the timeout, Guard G3 in the auxiliary state machine is activated only if the node is in state RECOVER at the time.

Main State Machine. The main state machine, which is given in Figure 7, is responsible for generating the pulse events and operates as follows. If a node is in state PULSE, it generates a local pulse event and sends a PULSE message to all other nodes. Now suppose a node $u \in G$ transitions to state PULSE. Two things can happen:

- If a node $u \in G$ is in state PULSE and observes at least $n - f$ nodes also generating a pulse within a short enough time window (Guard G1), it is possible that all correct nodes generated a pulse in a synchronised fashion. If this happens, then Guard G1 ensures that node u proceeds to the state WAIT. As the name suggests, the WAIT state is used to wait before generating a new pulse, ensuring that pulses obey the desired frequency bounds.
- Otherwise, if a node is certain that not all correct nodes are synchronised, it transitions from PULSE to state RECOVER (Guard G1’).

Once a node is in either WAIT or RECOVER, it will not leave the state before the consensus algorithm outputs “1”, as Guard G2 needs to be satisfied in order for a transition to PULSE to take place. The simulation of consensus is handled by the auxiliary state machine, which we discuss below. The nodes use consensus to agree whether sufficiently many nodes transitioned to the WAIT state within a small enough time window. If the system has stabilised, all correct nodes transition to WAIT almost synchronously, and hence, after stabilisation every correct node always uses input “1” for the consensus instance.

Once a node transitions to state WAIT, the node keeps track of how long it has been there. If the node observes that it has been there longer than it would take for a consensus simulation to complete under correct operation (indicating that the system has not yet stabilised), it transitions to state RECOVER. Also, if the consensus instance outputs “0”, the node knows something is wrong and transitions to RECOVER. During the stabilisation phase, nodes that transition to RECOVER refrain from using input “1” for any consensus routine before the local timeout T_{active} expires; we refer to the discussion of the auxiliary state machine.

Once the system stabilises, the behaviour of the main state machine is simple, as only Guard G1 and Guard G2 can be satisfied. This implies that correct nodes alternate between the PULSE and WAIT states. Under stabilised operation, we get that all correct nodes:

- transition to PULSE within a time window of length $2d$,
- observe that at least $n - f$ nodes transitioned to PULSE within a short enough time window ensuring that Guard G1 is satisfied at every correct node,
- transition to WAIT within a time window of length $O(d)$,
- correctly initialise a simulation of the consensus algorithm C with input “1”, as correct nodes transitioned to WAIT in a synchronised fashion (see auxiliary state machine),
- all correct nodes remain in WAIT until Guard G2 or Guard G2’ become satisfied.

Finally, we ensure that (after stabilisation) all correct nodes remain in state `WAIT` in the main state machine longer than it takes to properly initialise and simulate a consensus instance. This is achieved by using the T_{wait} timeout in Guard G2'. Due to the validity property of the consensus routine and the fact that all correct nodes use input 1, this entails that Guard G2 is always satisfied before Guard G2', such that all correct nodes again transition to `PULSE` within a time window of length $2d$.

Auxiliary State Machine. The auxiliary state machine given in Figure 8 is slightly more involved. However, the basic idea is simple:

- (a) nodes try to check whether at least $n - f$ nodes transition to the `WAIT` state *in a short enough time window* (that is, a time window consistent with correct operation) and
- (b) then use a consensus routine to agree whether all nodes saw this.

Assuming that all correct nodes participate in the simulation of consensus, we get the following:

- If the consensus algorithm `C` outputs “0”, then some correct node did not see $n - f$ nodes transitioning to `WAIT` in a short time window, and hence, the system has not yet stabilised.
- If the consensus algorithm `C` outputs “1”, then all correct nodes agree that a transition to `WAIT` happened recently.

In particular, the idea is that when the system operates correctly, the consensus simulation will always succeed and output “1” at every correct node.

The above idea is implemented in the auxiliary state machine as follows. Suppose that a correct node $u \in G$ is in the `LISTEN` state and the local timeout T_{active} is not about to expire (recall that T_{active} is only reset by the resynchronisation signal). Node u remains in this state until it is certain that at least one correct node transitions to `WAIT` in the main state machine. Once this happens, Guard G4 is satisfied and node u transitions to the `READ` state. In the `READ` state, node u waits for a while to see whether it observes (1) at least $n - f$ nodes transitioning to `WAIT` in a short time window or (2) less than $n - f$ nodes doing this.

In case (1), node u can be certain that at least $n - 2f > f$ correct nodes transitioned to `WAIT`. Thus, node u can also be certain that every correct node observes at least $f + 1$ correct nodes transitioning to `WAIT`; this will be a key property later. In case (2), node u can be certain that the system has not stabilised. If case (1) happens, we have that Guard G5 is eventually satisfied. Node u then transitions to `INPUT 1` indicating that node u is willing to use input “1” in the next simulation of consensus *unless* it is in the `RECOVER` state in the main state machine. In case (2), we get that Guard G5' becomes satisfied and u transitions to `INPUT 0`. This means that u insists on using input “0” for the next consensus simulation.

Once node $u \in G$ transitions to either `INPUT 0` or `INPUT 1`, it will remain there until the local timeout of length T_2 expires (see Guard G6, Guard G6', and Guard G7). However, if Guard G4 becomes satisfied *while* node u is in either of the input states, then the local timeout is reset again. We do this because, if Guard G4 becomes satisfied while u is in one of the input states, (i) the same may be true for other correct nodes that are in state `LISTEN` and (ii) node u can be certain that the system has not stabilised. Resetting the timeout helps in ensuring that all correct nodes jointly start the next consensus instance (guaranteeing correct simulation), if Guard G4 is satisfied at all correct nodes at roughly the same time. In case this does not happen, resetting the timeout at least makes sure that there will be a time when *no* correct node is currently trying to simulate a consensus instance. These properties are critical for our proof of stabilisation.

5.3 Outline of the Proof

The key difficulty in achieving stabilisation is to ensure the proper simulation of a consensus routine despite the arbitrary initial state. In particular, after the transient faults cease, we might have some nodes attempting to execute consensus, whereas some do not. Moreover, nodes that are simulating consensus might be simulating different rounds of the consensus routine, and so on. To show that such disarray cannot last indefinitely long, we use the following arguments:

- if some correct node attempts to use input “1” for consensus, then at least $f + 1$ correct nodes have transitioned to `WAIT` in the main state machine (Lemma 4), that is, all correct nodes see if some other correct node might be initialising a new consensus instance with input “1” soon,
- if some correct node transitions to `WAIT` at time t , then there is a long interval of length $\Theta(T_2)$ during which no correct node transitions to `WAIT` (Lemma 5), that is, correct nodes cannot transition to `WAIT` state too often,
- if some correct node attempts to use input “1” for consensus, then all correct nodes initialise a new consensus instance within a time window of length $\tau \in \Theta((1 - 1/\vartheta)T_2)$ (Lemma 6),
- if all correct nodes initialise a new consensus instance within a time window of length τ , then all correct nodes participate in the same consensus instance and successfully simulate an entire execution of `C` (Lemma 7).

The idea is that the timeout T_2 will be sufficiently large to ensure that consensus instances are well-separated: if a consensus instance is initialised with input “1” at some correct node, then there is enough time to properly simulate a complete execution of the consensus routine before any correct node attempts to start a new instance of consensus.

Once we have established the above properties, it is easy to see that if synchronisation is established, then it *persists*. More specifically, we argue that if all correct nodes transition to `PULSE` at most time $2d$ apart, then all correct nodes initialise a new consensus instance within a time window of length τ using input “1” (Lemma 8). Thus, the system stabilises if all correct nodes eventually generate a pulse with skew at most $2d$.

Accordingly, a substantial part of the proof is arguing that all nodes eventually transition to `PULSE` within time window of $2d$. To see that this is bound to occur eventually, we consider an interval $[\alpha, \beta]$ of length $\Theta(R)$ and use the following line of reasoning:

- if all correct nodes are simultaneously in state `RECOVER` at some time before timeout T_{active} expires at any correct node, then Guard G3 in the auxiliary state machine becomes satisfied at all correct nodes and a new consensus instance with all-1 input is initialised within a time window of length τ (Lemma 9),
- if some correct node attempts to use input “1” during the interval $[\alpha, \beta]$, then either (a) all correct nodes end up in `RECOVER` before timeout T_{active} expires at any node or (b) all correct nodes eventually transition to `PULSE` within time $2d$ (Lemma 10),
- if no correct node attempts to use input “1” during the time interval $[\alpha, \beta]$, all correct nodes will be in state `RECOVER` before the timeout T_{active} expires at any node (Lemma 14).

In either of the latter two cases, we can use the first argument to guarantee stabilisation (Corollary 4 and Corollary 5). Finally, we need to argue that all the timeouts employed in the construction can be set so that our arguments work out. The constraints related to all the timeouts are summarised in Table 3 and Lemma 16 shows that these can be satisfied. We now proceed to formalise and prove the above arguments in detail. The structure of the proof is summarised in Figure 9.

Table 3. The Timeout Conditions Employed in the Construction of Section 5

(5)	$d, \vartheta \in O(1)$
(6)	$T_1 = 3\vartheta d$
(7)	$T_{\text{listen}} = (\vartheta - 1)T_1 + 3\vartheta d$
(8)	$T_2 > \vartheta(T_{\text{listen}} + 3T_1 + 3d)$
(9)	$(2/\vartheta - 1)T_2 > 2T_{\text{listen}} + T_{\text{consensus}} + 5T_1 + 4d$
(10)	$\tau = \max \{ (1 - 1/\vartheta)T_2 + T_{\text{listen}} + d + \max\{T_{\text{listen}} + d, 3T_1 + 2d\}, (1 - 1/\vartheta)T_{\text{active}} + \rho \}$
(11)	$T_{\text{consensus}} = \vartheta(\tau + T(R))$
(12)	$T_{\text{wait}} = T_2 + T_{\text{consensus}}$
(13)	$T_{\text{active}} \geq 4T_2 + T_{\text{listen}} + \vartheta(T_{\text{listen}} + T_{\text{wait}} - 5T_1 - 4d + \rho)$
(14)	$T_{\text{active}} \geq 2T_2 + T_{\text{consensus}} + \vartheta(2T_{\text{listen}} + T_1 + T_{\text{wait}} + 3d + 2T_2 + 2T_{\text{consensus}})$

Fig. 9. The overall structure of the proof of Theorem 2. The bold rectangles denote results that are informally discussed in Section 5.3.

5.4 Analysing the State Machines

First let us observe that, after a short time, the arbitrary initial contents of the sliding window message buffers have been cleared.

Remark 2. By time $t = \max\{T_1, T_{\text{listen}}\} + d \in O(\vartheta^2 d)$ the sliding window memory buffers used in Guard G2, Guard G4, and Guard G5 for each $u \in G$ have valid contents: if the buffer of Guard G2 contains a message m from $v \in G$ at any time $t' \geq t$, then $v \in G$ sent the message m during $(t - T_1 - d, t)$; similarly, for Guard G4 and Guard G5, this holds for the interval $(t - T_{\text{listen}} - d, t)$.

Without loss of generality, we assume that this has happened by time 0. Moreover, we assume that every correct node received the resynchronisation signal during the time interval $[0, \rho)$. Thus,

the contents of message buffers are valid from time 0 on and every node has reset its T_{active} timeout during $[0, \rho)$. Hence, the timeout T_{active} expires at any node $u \in G$ during $[T_{\text{active}}/\vartheta, T_{\text{active}} + \rho)$.

We use $T(R) \in O(\vartheta^2 dR)$ to denote the maximum time a simulation of the R -round consensus routine **C** takes when employing the non-stabilising pulse synchronisation algorithm given in Section 4. We assume that the consensus routine **C** is silent, as by Theorem 5, we can convert any consensus routine into a silent one without any asymptotic loss in the running time.

First, we highlight some useful properties of the simulation scheme implemented in the auxiliary state machine.

Remark 3. If node $v \in G$ transitions to **RUN 0** or **RUN 1** at time t , then the following holds:

- node v remains in the respective state during $[t, t + \tau)$,
- node v does not execute the first round of **C** before time $t + \tau$, and
- node v leaves the respective state and halts the simulation of **C** before time $t + T_{\text{consensus}}$.

Now let us start by showing that if some node transitions to **INPUT 1** in the auxiliary state machine, then there is a group of at least $f + 1$ correct nodes that transition to **WAIT** in the main state machine in rough synchrony.

LEMMA 4. *Suppose node $v \in G$ transitions to **INPUT 1** at time $t \geq T_{\text{listen}} + d$. Then there is a time $t' \in (t - T_{\text{listen}} - d, t)$ and a set $A \subseteq G$ with $|A| \geq f + 1$ such that each $w \in A$ transitions to **WAIT** during $[t', t]$.*

PROOF. Since v transitions to **INPUT 1**, it must have observed at least $n - f$ distinct **WAIT** messages within time T_{listen} in order to satisfy Guard G5. As $f < n/3$, we have that at least $f + 1$ of these messages came from nodes $A \subseteq G$, where $|A| \geq f + 1$. The claim follows by choosing t' to be the minimal time during $(t - T_{\text{listen}} - d, t)$ at which some $w \in A$ transitioned to **WAIT**. $\square$

Next we show that if some correct node transitions to **WAIT**, then this is soon followed by a long time interval during which no correct node transitions to **WAIT**. Thus, transitions to **WAIT** are well separated.

LEMMA 5. *Suppose node $v \in G$ transitions to **WAIT** at time $t \leq (T_{\text{active}} - T_2)/\vartheta$. Then no $u \in G$ transitions to **WAIT** during $[t + 3T_1 + d, t + T_2/\vartheta - 2T_1 - d)$.*

PROOF. Since $v \in G$ transitioned to **WAIT** at time t , it must have seen at least $n - f$ nodes transitioning to **PULSE** during the time interval $(t - 2T_1, t)$. Since $f < n/3$, it follows that $n - 2f \geq f + 1$ of these messages are from correct nodes. Let us denote this set of nodes by $A \subseteq G$.

Consider any node $w \in A$. As node w transitioned to **PULSE** during $(t - 2T_1 - d, t)$, it transitions to state **RECOVER** or to state **WAIT** at time $t_w \in (t - 2T_1 - d, t + T_1)$. Either way, as it also transitioned to **LISTEN**, transitioning to **PULSE** again requires to satisfy Guard G3 or one of Guard G6, Guard G6', and Guard G7 while being in states **INPUT 0** or **INPUT 1**, respectively. By assumption, Guard G3 is not satisfied before time $T_{\text{active}}/\vartheta \geq t + T_2/\vartheta$, and the other options require a timeout of T_2 to expire, which takes at least time T_2/ϑ . It follows that w is not in state **PULSE** during $[t + T_1, t + T_2/\vartheta - 2T_1 - d)$.

We conclude that no $w \in A$ is observed transitioning to **PULSE** during $[t + T_1 + d, t + T_2/\vartheta - 2T_1 - d)$. Since $|A| > f$, we get that no $u \in G$ can activate Guard G1 and transition to **WAIT** during $[t + 3T_1 + d, t + T_2/\vartheta - 2T_1 - d)$, as $n - |A| < n - f$. $\square$

Using the previous lemmas, we can show that if some correct node transitions to state **INPUT 1** in the auxiliary state machine, then every correct node eventually initialises and participates in the same new consensus instance. That is, every correct node initialises the underlying Srikanth-Toueg pulse synchronisation algorithm within a time interval of length τ .

LEMMA 6. *Suppose node $u \in G$ transitions to INPUT 1 at time $t \in [T_{\text{listen}} + d, (T_{\text{active}} - T_2)/\vartheta]$. Then each $v \in G$ transitions to state RUN 0 or state RUN 1 at time $t_v \in [t_0, t_0 + \tau)$, where $t_0 = t - T_{\text{listen}} - d + T_2/\vartheta$. Moreover, Guard G4 cannot be satisfied at any node $v \in G$ during $[t + 3T_1 + 2d, t^* + T_1/\vartheta)$, where $t^* := \min_{v \in G} \{p(v, t + d)\}$.*

PROOF. By Lemma 4, there exists a set $A \subseteq G$ such that $|A| \geq f + 1$ and each $w \in A$ transitions to WAIT at a time $t_w \in (t - T_{\text{listen}} - d, t)$. This implies that Guard G4, and thus also Guard G9, becomes satisfied for $v \in G$ at time $t'_v \in [t - T_{\text{listen}} - d, t + d)$. Thus, every $v \in G$ transitions to state READ, INPUT 0, or INPUT 1 at time t'_v ; note that if v was in state INPUT 0 or INPUT 1 before this happened, it transitions back to the same state due to Guard G4 being activated and resets its local T_2 timer. Moreover, by time $t'_v \leq r_v < t'_v + T_{\text{listen}} < t + d + T_{\text{listen}}$ node v transitions to either INPUT 0 or INPUT 1, as either Guard G5 or Guard G5' becomes activated in case v transitions to state READ at time t'_v .

Now we have that node v remains in either INPUT 1 or INPUT 0 for the interval $[r_v, r_v + T_2/\vartheta)$, as none of Guard G6, Guard G6', and Guard G7 are satisfied before the local timer T_2 expires. Moreover, by applying Lemma 5 to any $w \in A$, we get that no $v \in G$ transitions to WAIT during the interval

$$[t_w + 3T_1 + d, t_w + T_2/\vartheta - 2T_1 - d] \supseteq (t + 3T_1 + d, t + T_2/\vartheta - 2T_1 - T_{\text{listen}} - 2d).$$

Recall that for each $v \in G$, $t'_v < t + d$. After this time, v cannot transition to PULSE again without transitioning to RUN 0 or RUN 1 first. Since $t + T_2/\vartheta - 2T_1 - T_{\text{listen}} - 2d > t + T_1 + d$ by Constraint (8), we get that every $w \in G$ has arrived in state WAIT or RECOVER by time $t + T_2/\vartheta - 2T_1 - T_{\text{listen}} - 2d$. Thus, no such node transitions to state WAIT during $[t + T_2/\vartheta - 2T_1 - T_{\text{listen}} - 2d, t^* + T_1/\vartheta)$: first, it must transition to PULSE, which requires to satisfy Guard G2, i.e., transitioning to state OUTPUT 1, and then a timeout of T_1 must expire; here, we use that we already observed that $t + T_2/\vartheta - 2T_1 - T_{\text{listen}} - 2d > t + d$, i.e., by definition the first node $w \in G$ to transition to PULSE after time $t + T_2/\vartheta - T_1 - T_{\text{listen}} - 2d$ does so at time t^* . We conclude that Guard G4 cannot be satisfied at any $v \in G$ during the interval $[t + 3T_1 + 2d, t^* + T_1/\vartheta)$, i.e., the second claim of the lemma holds.

We proceed to showing that each $v \in G$ transitions to state RUN 0 or state RUN 1 at time $t_v \in [t_0, t_0 + \tau)$, i.e., the first claim of the lemma. To this end, observe that w transitions to either state RUN 0 or RUN 1 at some time $t' \in (t'_w, t^*)$. By the above observations, $t' \geq r_w + T_2/\vartheta \geq t - T_{\text{listen}} - d + T_2/\vartheta = t_0$. Node w initialises the Srikant-Toueg algorithm given in Figure 5 locally at time t' . In particular, by the properties of the simulation algorithm given in Remark 3, we have that w waits until time $t' + \tau$ before starting the simulation of C , and hence, w remains in RUN 0 or RUN 1 at least until time $t' + \tau$ before the simulation of C produces an output value. Thus, we get that $t^* \geq t' + \tau \geq t_0 + \tau$.

Recall that each $v \in G$ resets timeout T_2 at time $r_v \in [t'_v, t + d + T_{\text{listen}}] \subseteq [t - T_{\text{listen}} - d, t + T_{\text{listen}} + d)$ and does not reset it during $[t + 3T_1 + 2d, t^* + T_1/\vartheta)$, as it does not satisfy Guard G4 at any time from this interval. When T_2 expires at v , it transitions to RUN 0 or RUN 1. Because $t^* + T_1/\vartheta > t_0 + \tau \geq t + \max\{T_{\text{listen}} + d, 3T_1 + 2d\} + T_2$ by Constraint (10), this happens at time $t_v \in [t - T_{\text{listen}} - d + T_2/\vartheta, t_0 + \tau] = [t_0, t_0 + \tau]$, as claimed. $\square$

Next, we show, that if all correct nodes initialise a new instance of the Srikant-Toueg pulse synchronisation algorithm within a time interval of length τ , then every correct node initialises, participates in, and successfully completes simulation of the consensus routine C .

LEMMA 7. *Suppose there exists a time t_0 such that each node $v \in G$ transitions to RUN 0 or RUN 1 at some time $t_v \in [t_0, t_0 + \tau)$. Let t' be the minimal time larger than t_0 at which some $u \in G$ transitions*

to either OUTPUT 0 or OUTPUT 1. If Guard G4 is not satisfied at any $v \in G$ during $[t_0, t' + 2d)$, then $t' \leq t_0 + T_{\text{consensus}}/\vartheta - 2d$ and there are times $t'_v \in [t', t' + 2d)$, $v \in G$, such that:

- each $v \in G$ transitions to OUTPUT 1 or OUTPUT 0 at time t'_v (termination),
- this state is the same for each $v \in G$ (agreement), and
- if each $v \in G$ transitioned to state RUN 1 at time t_v , then this state is OUTPUT 1 (validity).

PROOF. When $v \in G$ transitions to either RUN 0 or RUN 1, it sends an initialisation signal to the non-self-stabilising Srikanth-Toueg algorithm described in Section 4. Using the clock given by this algorithm, the nodes simulate the consensus algorithm C. If node $v \in G$ enters state RUN 0, it uses input “0” for C. Otherwise, if v enters RUN 1 it uses input “1”.

Note that Theorem 4 implies that, if all nodes initialise the Srikanth-Toueg algorithm within time τ apart, then the simulation of C takes at most $\tau + T(R) \in O(\vartheta^2 d(\tau + R))$ time. Moreover, all nodes will declare the output in the same round, and hence, declare the output within a time window of $2d$, as the skew of the pulses is at most $2d$.

Now let us consider the simulation taking place in the auxiliary state machine. If Guard G4 is not satisfied during $[t_0, t' + 2d)$ and the timer $T_{\text{consensus}}$ does not expire at any node $v \in G$ during the simulation, then by time $t' + 2d \leq t_0 + \tau + T(R) \in O(\vartheta^2 d(\tau + R))$ the nodes have simulated R rounds of C and declared an output. By assumption, Guard G4 cannot be satisfied prior to time $t' + 2d$. At node $v \in G$, the timer $T_{\text{consensus}}$ is reset at time $t_v \geq t_0$. Hence, it cannot expire again earlier than time $t_0 + T_{\text{consensus}}/\vartheta \geq t_0 + \tau + T(R)$ by Constraint (11). Hence, the simulation succeeds.

Since the simulation of the consensus routine C completes at each $v \in G$ at some time $t'_v \in [t', t' + 2d)$, we get that Guard G8 or Guard G9 is satisfied at time t'_v at node v . Hence, v transitions to either of the output states depending on the output value of C. The last two claims of the lemma follow from the agreement and validity properties of the consensus routine C. $\square$

Now we can show that if all correct nodes transition to PULSE within a time window of length $2d$, then all correct nodes remain synchronised with skew $2d$ and controlled accuracy bounds. Thus, the system stabilises.

LEMMA 8. *Suppose there exists an interval $[t, t + 2d)$ such that for all $v \in G$ it holds that $p(v, t_v) = 1$ for some $t_v \in [t, t + 2d)$. Then there exists $t' \in [t + T_2/\vartheta, t + (T_2 + T_{\text{consensus}})/\vartheta - 2d)$ such that $p_{\text{next}}(v, t_v) \in [t', t' + 2d)$ for all $v \in G$.*

PROOF. First, observe that if any node $v \in G$ transitions to PULSE at time t_v , then node v transitions to state LISTEN in the auxiliary state machine at time t_v . To see this, note that node v must have transitioned to OUTPUT 1 in the auxiliary state machine at time t_v in order to satisfy Guard G2 leading to state PULSE. Furthermore, once this happens, node v transitions immediately from OUTPUT 1 to LISTEN in the auxiliary state machine. Note that no $v \in G$ transitioned to WAIT during $(t_v - T_2/\vartheta, t_v)$, as v waits for at least this time before initiating another consensus instance after transitioning to OUTPUT 1. Hence, Constraint (8) ensures that no correct node stores any WAIT messages from any other correct node in its WAIT sliding window buffer. It follows that each $v \in G$ is in state LISTEN at time t_v and cannot leave before time $t + T_1/\vartheta$.

Next, note that $v \in G$ will not transition to RECOVER before time $t_v + T_1/\vartheta \geq t + 3d$ by Guard G1 and Constraint (6). By assumption, every $u \in G$ transitions to PULSE by time $t + 2d$, and thus, node v observes a PULSE message from at least $n - f$ correct nodes $u \in G$ by time $t + 3d$. Thus, all correct nodes observe a PULSE message from at least $n - f$ nodes during $[t, t + T_1/\vartheta) = [t, t + 3d)$ satisfying Guard G1. Hence, every correct node $v \in G$ transitions to WAIT during the interval $[t + T_1/\vartheta, t + T_1 + 2d)$ and remains there until Guard G2 or Guard G2' is activated. Denote by t'_v the next transition of $v \in G$ to OUTPUT 1 or OUTPUT 0 after time t_v ($t'_v := \infty$ if no such time exists) and set $t' := \min_{v \in G} \{t'_v\}$. Guard G2' cannot be activated before time $\min\{t', t + T_{\text{wait}}/\vartheta\}$.

Now let us consider the auxiliary state machine. From the above reasoning, we get that every node $v \in G$ will observe at least $n - f$ nodes transitioning from PULSE to WAIT during the interval $[t + T_1/\vartheta, t + T_1 + 3d]$. As we have $T_{\text{listen}}/\vartheta \geq (1 - 1/\vartheta)T_1 + 3d$ by Constraint (7), both Guard G4 and Guard G5 become satisfied for v during the same interval. Thus, node v transitions to state INPUT 1 during the interval $[t + T_1/\vartheta, t + T_1 + 3d]$. Here, we use that Guard G3 is not active at any $v \in G$ that is not in RECOVER, implying that $t' \geq \min\{t + T_2/\vartheta, t + T_{\text{wait}}/\vartheta\} = t + T_2/\vartheta > t + T_1 + 3d$ by Constraint (8) and Constraint (12). Thus, Guard G2' cannot become active before time $t + T_1 + 3d$, yielding that each $v \in G$ transitions to READ before Guard G3 can become active. We claim that v transitions to RUN 1 during $[t + (T_1 + T_2)/\vartheta, t + T_1 + T_2 + 3d] \subseteq [t + T_2/\vartheta, t + T_2/\vartheta + \tau]$ by Constraint (10). Assuming otherwise, some $v \in G$ would have to transition to state WAIT before time $T_2/\vartheta + \tau$. However, $t' \geq T_2/\vartheta + \tau$ by Remark 3 and $t + T_{\text{wait}}/\vartheta = t + (T_2 + T_{\text{consensus}})/\vartheta > t + T_2/\vartheta + \tau$ by Constraint (12) and Constraint (11).

Note that $t + T_1 + 4d < t + T_2/\vartheta$ by Constraint (6) and Constraint (8) and that no correct node transitions to WAIT again after time $t + T_1 + 3d$ before transitioning to OUTPUT 1 and spending, by Constraint (6), at least $T_1/\vartheta > 2d$ time in PULSE. Therefore, we can apply Lemma 7 with $t_0 = t + T_2/\vartheta$, yielding a time $t' < t + T_2/\vartheta + T_{\text{consensus}}/\vartheta - 2d$ such that each $v \in G$ transitions to OUTPUT 1 in the auxiliary state machine at time $t'_v \in [t', t' + 2d]$. This triggers a transition to PULSE in the main state machine. $\square$

5.5 Ensuring Stabilisation

We showed above that if all correct nodes eventually generate a pulse with skew $2d$, then the pulse synchronisation routine stabilises. In this section, we show that this is bound to happen. The idea is that if stabilisation does not take place within a certain interval, then all nodes end up being simultaneously in state RECOVER in the main state machine and state LISTEN in the auxiliary state machine, until eventually timeout T_{active} expires. This in turn allows the “passive” stabilisation mechanism to activate by having timer T_{active} expire at every node and stabilise the system as shown by the lemma below.

LEMMA 9. *Let $t < T_{\text{active}}/\vartheta$ and $t^* = \min_{v \in G} \{p_{\text{next}}(v, t)\}$. Suppose the following holds for every node $v \in G$:*

- *node v is in state RECOVER and LISTEN at time t , and*
- *Guard G4 is not satisfied at v during $[t, t^* + 2d]$,*

Then $t^ < T_{\text{active}} + \rho + T_{\text{consensus}}/\vartheta$ and every $v \in G$ transitions to PULSE at time $t_v \in [t^*, t^* + 2d]$.*

PROOF. Observe that Guard G3 is not satisfied before time $T_{\text{active}}/\vartheta$. As Guard G4 is not satisfied during $[t, t^*]$, no correct node leaves state T_{listen} before time $T_{\text{active}}/\vartheta$. Since no $v \in G$ can activate Guard G2 during $[t, t^*]$, we have that every $v \in G$ remains in state RECOVER during this interval. Let $t_0 \in [T_{\text{active}}/\vartheta, T_{\text{active}} + \rho)$ be the minimal time (after time t) when Guard G3 becomes satisfied at some node $v \in G$. From the properties of the simulation algorithm given in Remark 3, we get that no $w \in G$ transitions away from RUN 1 before time $t_0 + \tau \geq T_{\text{active}}/\vartheta + \tau$.

Since $\tau \geq (1 - 1/\vartheta)T_{\text{active}} + \rho$ by Constraint (10), we conclude that no $w \in G$ transitions to OUTPUT 1 (and thus PULSE) before time $T_{\text{active}} + \rho$. Therefore, each $w \in G$ transitions to RUN 1 at some time $r_w \in [T_{\text{active}}/\vartheta, T_{\text{active}} + \rho] \subseteq [t_0, t_0 + \tau)$. Recall that Guard G4 is not satisfied during $[t_0, t^* + 2d]$. Hence, we can apply Lemma 7 to the interval $[t_0, t_0 + \tau)$, implying that every $w \in G$ transitions to OUTPUT 1 at some time $t_w \in [t^*, t^* + 2d]$. Thus, each $w \in G$ transitions from RECOVER to PULSE at time t_w . Finally, Lemma 7 also guarantees that $t^* < t_0 + T_{\text{consensus}}/\vartheta \leq T_{\text{active}} + \rho + T_{\text{consensus}}/\vartheta$. $\square$

We will now establish a series of lemmas in order to show that we can either apply Lemma 8 directly or in conjunction with Lemma 9 to guarantee stabilisation. In the remainder of this section, we define the following abbreviations:

$$\begin{aligned}\alpha &= T_{\text{listen}} + d \\ \beta &= \alpha + T_{\text{wait}} + \gamma + 4T_1 + 3d + 2\delta \\ \beta' &= (T_{\text{active}} - T_2 - T_{\text{consensus}})/\vartheta \\ \delta &= T_1 + d + 2T_{\text{listen}} + T_2 + T_{\text{consensus}} \\ \gamma &= T_2/\vartheta - T_{\text{listen}} - 5T_1 - 3d.\end{aligned}$$

Remark 4. We have that $\gamma < \delta < \beta \leq \beta' < T_{\text{active}}/\vartheta$, where the inequality $\beta \leq \beta'$ is equivalent to Constraint (14).

In the following, we consider the time intervals $[\alpha, \beta]$ and $[\alpha, \beta']$ that depend on different timeouts. We distinguish between two cases and show that in either case stabilisation is ensured. The cases are:

- no correct node transitions to INPUT 1 during $[\alpha, \beta]$, and
- some correct node transitions to INPUT 1 during $[\alpha, \beta']$.

These cases correspond to our proof strategy described in Section 5.3: If the first case occurs, all nodes end up in the RECOVER state and the “passive” stabilisation mechanism guarantees stabilisation after the timer T_{active} expires. If the first case does not hold, then we must be in the latter case if $\beta' \geq \beta$ holds, which happens to be true when Constraint (14) is satisfied. In the second case, we show that a consensus instance will be run by all correct processors, which then can be used to ensure that all nodes agree that a pulse should be generated (output “1”) or that the system has not stabilised (output “0”), which leads to everyone transitioning to the RECOVER state. We start by analysing the case where some correct node transitions to INPUT 1 during the interval $[\alpha, \beta']$.

LEMMA 10. *Suppose node $v \in G$ transitions to INPUT 1 at time $t \in [\alpha, \beta']$. Then there exists a time $t' \in [t, T_{\text{active}}/\vartheta - 2d)$ such that one of the following holds:*

- (1) *every $u \in G$ satisfies $p(u, t_u) = 1$ for some $t_u \in [t', t' + 2d)$, or*
- (2) *every $u \in G$ is in state RECOVER and LISTEN at time $t' + 2d$ and Guard G4 is not satisfied at u during $[t' + 2d, t^* + 2d]$, where $t^* := \min_{w \in G} \{p(w, t' + 2d)\}$.*

PROOF. By Constraint (6), we have $T_1/\vartheta > 2d$. As $\alpha \leq t \leq \beta' < (T_{\text{active}} - T_2)/\vartheta$, we can apply Lemma 6 to time t . Due to Constraint (8), we can apply Lemma 7 with time $t_0 = t - T_{\text{listen}} - d + T_2/\vartheta$, yielding a time $t' \leq \beta' - T_{\text{listen}} - 3d + T_2/\vartheta + T_{\text{consensus}}/\vartheta < T_{\text{active}}/\vartheta - 2d$ such that each $u \in G$ transitions to the same output state OUTPUT 0 or OUTPUT 1 in the auxiliary state machine at time $t_u \in [t', t' + 2d)$. If this state is OUTPUT 1, each $u \in G$ transitions to OUTPUT 1 at time $t_u \in [t', t' + 2d)$. This implies that Guard G2 is activated and u transitions to PULSE so that $p(u, t_u) = 1$. If this state is OUTPUT 0, Guard G2' implies that each $u \in G$ either remains in or transitions to state RECOVER at time t_u . Moreover, node u immediately transitions to state LISTEN in the auxiliary state machine. Finally, note that Lemma 6 also states that Guard G4 cannot be satisfied again before time $t^* + T_1/\vartheta > t^* + 2d$. $\square$

COROLLARY 4. *Suppose node $v \in G$ transitions to INPUT 1 at time $t \in [\alpha, \beta']$. Then there exists $t' < T_{\text{active}} + \rho + T_{\text{consensus}}/\vartheta$ such that every $u \in G$ satisfies $p(v, t_u) = 1$ for some $t_u \in [t', t' + 2d)$.*

PROOF. We apply Lemma 10. In the first case of Lemma 10, the claim immediately follows. In the second case, it follows by applying Lemma 9. $\square$

We now turn our attention to the other case, where no node $v \in G$ transitions to INPUT 1 during the time interval $[\alpha, \beta]$.

LEMMA 11. *If no $v \in G$ transitions to INPUT 1 during $[\alpha, \beta]$, then no $v \in G$ transitions to state RUN 1 during $[\alpha + T_1 + T_{\text{wait}}, \beta]$.*

PROOF. Observe that any $v \in G$ that does not transition to PULSE during $[\alpha, \alpha + T_1 + T_{\text{wait}}]$ must transition to RECOVER at some time from that interval once Guard G2' is satisfied. Note that leaving state RECOVER requires Guard G2 to be satisfied, and hence, a transition to OUTPUT 1 in the auxiliary state machine. However, as no node $v \in G$ transitions to INPUT 1 during $[\alpha, \beta]$, it follows that each $v \in G$ that is in state INPUT 1 in the auxiliary state machine at time $t \in [\alpha + T_1 + T_{\text{wait}}, \beta]$ is also in state RECOVER in the main state machine. Thus, Guard G6 cannot be satisfied. We conclude that no correct node transitions to state RUN 1 during the interval $[\alpha + T_1 + T_{\text{wait}}, \beta]$. $\square$

We now show that if Guard G4 cannot be satisfied for some time, then there exists an interval during which no correct node is simulating a consensus instance.

LEMMA 12. *Let $t \in (d, \beta - 2\delta)$ and suppose Guard G4 is not satisfied during the interval $[t, t + \gamma]$. Then there exists a time $t' \in [t, \beta - \delta]$ such that no $v \in G$ is in state RUN 0 or RUN 1 during $(t' - d, t']$.*

PROOF. Observe that Guard G3 cannot be satisfied before time $t + \gamma + T_2/\vartheta < T_{\text{active}}/\vartheta$ at any correct node, and by the assumption, Guard G4 is not satisfied during the interval $[t, t + \gamma]$. We proceed via a case analysis.

First, suppose $v \in G$ is in state LISTEN at time t . As Guard G3 or Guard G4 cannot be satisfied during $[t, t + \gamma]$ node, v remains in LISTEN until time $t + \gamma$. Moreover, if v leaves LISTEN during $[t + \gamma, t + \gamma + T_2/\vartheta]$, it must do so by transitioning to READ. Hence, it cannot transition to RUN 0 or RUN 1 before Guard G6, Guard G6', or Guard G7 is satisfied. Either way, v cannot reach states RUN 0 or RUN 1 before time $t + \gamma + T_2/\vartheta$. Hence, in this case, v is not in the two execution states during $I_1 = [t, t + \gamma + T_2/\vartheta] = [t, t + 2T_2/\vartheta - T_{\text{listen}} - 5T_1 - 3d]$.

Let us consider the second case, where v is not in LISTEN at time t . Note that the timer T_2 cannot be reset at v during the interval $[t + T_{\text{listen}}, t + \gamma]$, as this can happen only if v transitions to INPUT 0 or INPUT 1 and Guard G4 cannot be satisfied during $[t, t + \gamma]$. Hence, the only way for this to happen is if v transitions from READ to either INPUT 0 or INPUT 1 during the interval $[t, t + T_{\text{listen}}]$.

It follows that v cannot transition to RUN 0 or RUN 1 during the interval $(t + T_{\text{listen}} + T_2, t + \gamma + T_2/\vartheta)$. Moreover, if v transitions to these states before $t + T_{\text{listen}} + T_2$, then v must transition away from them by time $t + T_{\text{listen}} + T_2 + T_{\text{consensus}}$, as Guard G8 or Guard G9 become satisfied. Therefore, node v is in LISTEN, READ, INPUT 0, or INPUT 1 during $I_2 = [t + T_{\text{listen}} + T_2 + T_{\text{consensus}}, t + \gamma + T_2/\vartheta]$. Applying Constraint (9), we get that

$$T_{\text{listen}} + T_2 + T_{\text{consensus}} < 2T_2/\vartheta - T_{\text{listen}} - 5T_1 - 4d,$$

and hence, by setting $t' = t + T_{\text{listen}} + T_2 + T_{\text{consensus}} + d < \beta - \delta$, we get that $(t' - d, t'] \subseteq I_1 \cap I_2$. That is, in either case, v is not in state RUN 0 or RUN 1 during this short interval. $\square$

Next, we show that the precondition of the previous lemma will be satisfied in due time.

LEMMA 13. *There exists $t \in [\alpha + T_1 + T_{\text{wait}} + d, \alpha + 4T_1 + T_{\text{wait}} + 3d + \gamma) \subset (d, \beta - 2\delta)$ such that Guard G4 is not satisfied during $[t, t + \gamma]$ at any $v \in G$.*

PROOF. Let $t' \in [\alpha + T_1 + T_{\text{wait}}, \alpha + T_1 + T_{\text{wait}} + d + \gamma]$ be a time when some $v \in G$ transitions to state WAIT. If such t' does not exist, the claim trivially holds for $t = \alpha + T_1 + T_{\text{wait}} + d$. Otherwise,

since $t' \leq T_{\text{listen}} + T_1 + T_{\text{wait}} + 2d + \gamma < (T_{\text{active}} - T_2)/\vartheta$ by Constraint (13), we can apply Lemma 5 to time t' , which yields that no $u \in G$ transitions to WAIT during

$$[t' + 3T_1 + d, t' + T_2/\vartheta - 2T_1 - d) = [t - T_{\text{listen}} - d, t + \gamma),$$

where $t = t' + T_{\text{listen}} + 3T_1 + 2d$. Thus, Guard G4 is not satisfied at any $v \in G$ during $[t, t + \gamma]$. $\square$

Using the above statements, we can now infer that if no node transitions to INPUT 1 for a while, this implies that no node transitions to PULSE for a certain period. Subsequently, we will use this to conclude that then the preconditions of Lemma 9 are satisfied.

LEMMA 14. *Suppose no $v \in G$ transitions to INPUT 1 during $[\alpha, \beta]$. Then no $v \in G$ transitions to PULSE during $[\beta - \delta, \beta]$.*

PROOF. First, we apply Lemma 13 to obtain an interval $[t, t + \gamma] \subseteq [\alpha + T_1 + T_{\text{wait}} + d, \beta - 2\delta]$ during which Guard G4 cannot be satisfied at any $v \in G$. Applying Lemma 12 to this interval yields an interval $(t' - d, t') \subset (t, \beta - \delta]$ during which no $v \in G$ is in state RUN 0 or RUN 1. This implies that no node is running a consensus instance during this time interval, and moreover, no messages from prior consensus instances are in transit to or arrive at any correct node at time t' . In particular, any $v \in G$ that (attempts to) simulate a consensus instance at time t' or later must first reinitialise the simulation by transitioning to RUN 0 or RUN 1.

Next, let us apply Lemma 11, to see that no v transitions to RUN 1 during the interval $[\alpha + T_1 + T_{\text{wait}}, \beta] \supset [t', \beta]$. Thus, if any node $v \in G$ attempts to simulate C, it must start the simulation by transitioning to RUN 0. This entails that any correct node attempting to simulate C does so with input 0. Because C is a silent consensus routine (see Definition 2), this entails that v does not send any message related to C unless it receives one from a correct node first, and in absence of such a message, it will not terminate with output 1. We conclude that no $v \in G$ transitions to OUTPUT 1 during $[t', \beta] \supseteq [\beta - \delta, \beta]$. The claim follows by observing that a transition to PULSE in the main state machine requires a transition to OUTPUT 1 in the auxiliary state machine. $\square$

LEMMA 15. *Suppose no $v \in G$ transitions to PULSE during $[\beta - \delta, \beta]$. Then, at time β , every $v \in G$ is in state RECOVER in the main state machine and state LISTEN in the auxiliary state machine. Moreover, Guard G4 is not satisfied during $[\beta, t^* + 2d)$, where $t^* = \min_{v \in G} \{p_{\text{next}}(v, \beta)\}$.*

PROOF. As no $v \in G$ transitions PULSE during $[\beta - \delta, \beta]$, either Guard G1 or Guard G2' lead each $v \in G$ to RECOVER. More precisely, every $v \in G$ is in state RECOVER of the main state machine during $[\beta - \delta + T_1 + T_{\text{wait}}, \beta]$. Next, observe that Guard G4 is not satisfied during $[\beta - \delta + T_1 + T_{\text{listen}} + d, \beta]$, and because $\beta < T_{\text{active}}/\vartheta$, Guard G3 cannot be active at any time from this interval either. It follows that each $v \in G$ is in state LISTEN of the auxiliary state machine during $[\beta - \delta + T_1 + d + 2T_{\text{listen}} + T_2 + T_{\text{consensus}}, \beta] = [\beta, \beta]$, i.e., the first claim of the lemma holds. For the second claim, observe that Guard G4 cannot be satisfied before the next transition to WAIT by a correct node occurs. This cannot happen before T_1/ϑ time has passed after a correct node transitioned to PULSE by Guard G1. Since $T_1/\vartheta > 2d$ by Constraint (6), the claim follows. $\square$

We can now show that if no INPUT 1 transitions occur during $[\alpha, \beta]$, then all nodes end up in the RECOVER state in the main state machine before T_{active} timeout expires at any node. This will eventually activate Guard G3 at every correct node, leading to a correct simulation of C with all 1 inputs.

COROLLARY 5. *Suppose no $v \in G$ transitions to INPUT 1 during $[\alpha, \beta]$. Then there exists a time $t < T_{\text{active}} + \rho + T_{\text{consensus}}$ such that every $v \in G$ transitions to PULSE at time $t_v \in [t, t + 2d)$.*

PROOF. By Lemma 14 and Lemma 15, the prerequisites of Lemma 9 are satisfied at time $t = \beta < T_{\text{active}}/\vartheta$. $\square$

It remains to show that the constraints in Table 3 can be satisfied for some $\vartheta > 1$ such that all timeouts are in $O(R)$.

LEMMA 16. *Let $1 < \vartheta < (2 + \sqrt{32})/7 \approx 1.094$. The constraints in Table 3 can be satisfied with all timeouts in $O(R)$.*

PROOF. Recall that R is the number of rounds the consensus routine needs to declare output and $T(R)$ is the time required to simulate the consensus routine. We parametrise T_2 and T_{active} using X and Y , where we require that $X \in \Theta(R)$ is chosen so that $T(R)/X \leq \varepsilon$, for a constant $\varepsilon > 0$ to be determined later. We then can set

$$\begin{aligned} T_1 &:= 3\vartheta d \\ T_{\text{listen}} &:= 3\vartheta^2 d \\ T_2 &:= X \\ \tau &:= \max\{(1 - 1/\vartheta)X + (3\vartheta^2 + 9\vartheta + 3)d, (1 - 1/\vartheta)Y + \rho\} \\ T_{\text{consensus}} &:= \vartheta(\tau + \varepsilon X) \\ T_{\text{wait}} &:= X + T_{\text{consensus}} \\ T_{\text{active}} &:= Y, \end{aligned}$$

immediately satisfying Constraint (6), Constraint (7), Constraint (10), Constraint (11), and Constraint (12). Moreover, Constraint (8) holds by requiring that X is at least a sufficiently large constant.

For the remaining constraints, denote by $C(\vartheta, d)$ a sufficiently large constant subsuming all terms that depend only on ϑ and d and abbreviate $T'_{\text{consensus}} := (\vartheta - 1) \max\{X, Y\} + \varepsilon\vartheta X$ (i.e., the non-constant terms of $T_{\text{consensus}}$). To satisfy Constraint (9), Constraint (13), and Constraint (14), it is then sufficient to guarantee that

$$\begin{aligned} (2/\vartheta - 1)X &> T'_{\text{consensus}} + C(\vartheta, d) \\ Y &\geq 4X + \vartheta(X + T'_{\text{consensus}}) + C(\vartheta, d) \\ Y &\geq 2X + T'_{\text{consensus}} + 3\vartheta(X + T'_{\text{consensus}}) + C(\vartheta, d), \end{aligned}$$

where the second inequality automatically holds if the third is satisfied. We also note that $Y \geq X$ is a necessary condition to satisfy the third inequality, implying that we may assume that $T'_{\text{consensus}} = (\vartheta - 1)Y + \varepsilon X$. We rearrange the remaining inequalities, yielding

$$\begin{aligned} (2/\vartheta - 1 - \varepsilon)X &> (\vartheta - 1)Y + C(\vartheta, d) \\ (2 + 2\vartheta - 3\vartheta^2)Y &\geq (2 + 3\vartheta(1 + \varepsilon\vartheta))X + C(\vartheta, d). \end{aligned} \tag{1}$$

Recall that ϑ and $C(\vartheta, d)$ are constant, and ε is a constant under our control. Hence, these inequalities can be satisfied if and only if

$$(2/\vartheta - 1)(2 + 2\vartheta - 3\vartheta^2) > (2 + 3\vartheta)(\vartheta - 1).$$

The above inequality holds for all $\vartheta \in (1, (2 + \sqrt{32})/7)$. Note that, as ϑ , ε , and $C(\vartheta, d)$ are constants, we can choose $X \in \Theta(R)$ as initially stated, implying that all timeouts are in $O(R)$, as desired. $\square$

COROLLARY 6. *For $\vartheta < (2 + \sqrt{32})/7$ and suitably chosen timeouts, in any execution there exists $t \in O(R)$ such that every $v \in G$ transitions to PULSE during the time interval $[t, t + 2d)$.*

PROOF. We choose the timeouts in accordance with Lemma 16. If no $v \in G$ transitions to INPUT 1 during $[\alpha, \beta]$, Corollary (5) yields the claim. Otherwise, some node $v \in G$ transitions to state INPUT 1 during the interval $[\alpha, \beta] \subseteq [\alpha, \beta']$ and the claim holds by Corollary 4. $\square$

Next, we observe that the accuracy bounds can be set to be within a constant factor apart from each other.

COROLLARY 7. *Let $\vartheta < (2 + \sqrt{32})/7$ and $\varphi_0(\vartheta) = 1 + 5(\vartheta - 1)/(2 + 2\vartheta - 3\vartheta^2) \in 1 + O(\vartheta - 1)$. For any constant $\varphi > \varphi_0(\vartheta)$, we can obtain accuracy bounds satisfying $\Phi^+/\Phi^- \leq \varphi$ and $\Phi^-, \Phi^+ \in \Theta(R)$.*

PROOF. By Lemma 8, the accuracy bounds we get from the constuction are $\Phi^- = T_2/\vartheta$ and $\Phi^+ = (T_2 + T_{\text{consensus}})/\vartheta$. Choosing the timeouts as in the proof of Lemma 16, we get that $\Phi^+/\Phi^- = 1 + (\vartheta - 1)Y/X + \varepsilon$, where ε is an arbitrarily small constant. Checking Inequality (1), we see that we can choose $Y/X = (2 + 3(1 + \varepsilon))/(2 + 2\vartheta - 3\vartheta^2)$. Choosing ε sufficiently small, the claim follows. $\square$

5.6 Proof of Theorem 2

Finally, we are ready to prove the main theorem of this section.

THEOREM 2. *Let $f \geq 0$, $n > 3f$, and $(2 + \sqrt{32})/7 > \vartheta > 1$. Suppose for a network of n nodes, there exists*

- *an f -resilient synchronous consensus algorithm **C**, and*
- *an f -resilient resynchronisation algorithm **B** with skew $\rho \in O(d)$ and sufficiently large separation window $\Psi \in O(R)$ that tolerates clock drift of ϑ ,*

*where **C** runs in $R = R(f)$ rounds and lets nodes send at most $M = M(f)$ bits per round and channel. Then there exists $\varphi_0(\vartheta) \in 1 + O(\vartheta - 1)$ such that, for any constant $\varphi > \varphi_0(\vartheta)$ and sufficiently large $T \in O(R)$, there exists an f -resilient pulse synchronisation algorithm **A** for n nodes that*

- *has skew $\sigma = 2d$,*
- *satisfies the accuracy bounds $\Phi^- = T$ and $\Phi^+ = T\varphi$,*
- *stabilises in $T(\mathbf{B}) + O(R)$ time, and*
- *has nodes send $M(\mathbf{B}) + O(M)$ bits per time unit and channel.*

PROOF. By the properties of the resynchronisation algorithm **B**, we get that a good resynchronisation pulse occurs within time $T(\mathbf{B})$. Once this happens, Corollary (6) shows all correct nodes transition to PULSE during $[t, t + 2d)$ for $t \in T(\mathbf{B}) + O(R)$. By Lemma 8, we get that the algorithm stabilises by time t and has skew $\sigma = 2d$. From Corollary (7), we get that the accuracy bounds can be set to be within factor φ apart without affecting the stabilisation time asymptotically.

To analyse the number of bits sent per time unit, first observe that the main state machine communicates whether a node transitions to PULSE or WAIT. This can be encoded using messages of size $O(1)$. Moreover, as node remains $\Omega(d)$ time in PULSE or WAIT, the main state machine sends only $O(1)$ bits per time unit. Second, the auxiliary state machine does not communicate apart from messages related to the simulation of consensus. The non-self-stabilising pulse synchronisation algorithm sends messages only when a node generates a pulse and the time between pulses is $\Omega(d)$. Thus, while simulating **C**, each node broadcasts at most $M(\mathbf{C}) + O(1)$ bits per time unit. $\square$

6 RESYNCHRONISATION ALGORITHMS

In this section, we give the second key component in the proof of Theorem 1. We show that given pulse synchronisation algorithms for networks of small size and with low resilience, it is possible to obtain resynchronisation algorithms for large networks and with high resilience. More precisely, we show the following theorem:

THEOREM 3. *Let $f, n_0, n_1 \in \mathbb{N}$ and $1 < \vartheta \leq 1.004$. Define*

$$n = n_0 + n_1, \quad f_0 = \lfloor (f - 1)/2 \rfloor, \quad f_1 = \lceil (f - 1)/2 \rceil.$$

For any $\Psi \in \Omega(1)$ and sufficiently small constant $\varphi > \varphi_0(\vartheta)$, there exists a bound $T_0 \in \Theta(\Psi)$ such that the following claim holds. If, for both $i \in \{0, 1\}$, there exists pulse synchronisation algorithm $\mathbf{A}_i$ that

- runs on n_i nodes and has resilience f_i ,
- has skew $\sigma = 2d$, and
- has accuracy bounds $\Phi_i^- = T$ and $\Phi_i^+ = T\varphi$, where $T_0 \leq T$ and $T \in O(\Psi)$,

then there exists a resynchronisation algorithm $\mathbf{B}$ that

- runs on n nodes and has resilience f ,
- has skew $\rho \in O(d)$ and separation window of length Ψ ,
- generates a resynchronisation pulse by time $T(\mathbf{B}) \in \max\{T(\mathbf{A}_0), T(\mathbf{A}_1)\} + O(\Psi)$, and
- has nodes send $M(\mathbf{B}) \in \max\{M(\mathbf{A}_0), M(\mathbf{A}_1)\} + O(1)$ bits per time unit and channel.

6.1 The High-Level Idea

Our goal is to devise a self-stabilising resynchronisation algorithm with skew $\rho \in O(d)$ and separation window Ψ for n nodes that tolerates $f < n/3$ faulty nodes. That is, we want an algorithm that guarantees that there exists a time t such that all correct nodes locally generate a single resynchronisation pulse during the interval $[t, t + \rho)$ and no new pulse during the interval $[t + \rho, t + \rho + \Psi)$. Note that a correct resynchronisation algorithm is also allowed to generate various kinds of spurious resynchronisation pulses, such as pulses that are followed by a new resynchronisation pulse too soon (i.e., before Ψ time units have passed) or pulses that are only generated by a subset of the correct nodes.

The Algorithm Idea. In order to illustrate the idea behind our resynchronisation algorithm, let us ignore clock drift and suppose we have two sources of pulses that generate pulses with fixed frequencies. Whenever either source generates a pulse, then a resynchronisation pulse is triggered as well. If the sources generate pulses with frequencies that are coprime multiples of (a sufficiently large) $C \in \Theta(\Psi)$, then we are guaranteed that eventually one of the sources produces a pulse followed by at least Ψ time units before a new pulse is generated by either of the two sources. See Figure 10(a) for an illustration.

Put otherwise, suppose $p_h(v, t) \in \{0, 1\}$ indicates whether node v observes the pulse source $h \in \{0, 1\}$ generating a pulse at time t . Using the above scheme, the output variable for the resynchronisation algorithm would be $r(v, t) = \max_{h \in \{0, 1\}} \{p_h(v, t)\}$. If, eventually, each source h generates a pulse roughly every C_h time units, setting $C_0 < C_1$ to be coprime integer multiples of $C \in \Theta(\Psi)$ (we allow for a constant-factor slack to deal with clock drift, etc.), we eventually have a time when a pulse is generated by one source, but no source will generate another pulse for at least Ψ time units.

Obviously, if we had such reliable self-stabilising pulse sources for n nodes and $f < n/3$ faulty nodes, then we would have effectively solved the pulse synchronisation problem already. However, our construction given in Section 5 relies on having a resynchronisation algorithm. Thus, in order to avoid this chicken-and-egg problem, we partition the set of n nodes into two and have each part run an instance of a pulse synchronisation algorithm with resilience almost $f/2$. That is, we take two *pulse synchronisation algorithms* with low resilience, and use these to obtain a *resynchronisation algorithm* with high resilience. This allows us to recursively construct resynchronisation algorithms starting from trivial pulse synchronisation algorithms that do not tolerate any faulty nodes.

The final obstacle to this construction is that we cannot guarantee that *both* instances with smaller resilience operate correctly, as the total number of faults exceeds the number that can be tolerated by each individual instance. We overcome this by enlisting the help of *all* nodes to check, for each instance, whether its output appears to satisfy the desired frequency bounds. If not, its

Fig. 10. Idea of the resynchronisation algorithm. We take two pulse sources (top two rows) with coprime frequencies and output the logical OR of the two sources (bottom row). Here, the dim gray lines delimit time intervals of length $C \in \Theta(\Psi)$. The blue regions indicate intervals when a correctly operating source should generate a pulse. In this example, the pulses of the first source should occur approximately every $2C$ time units, whereas the pulses of the second source should occur approximately every $3C$ time units. (a) Two correct sources that pulse with set frequencies. All correct nodes observe the same input pulses, and hence, produce the same output pulses. A good resynchronisation pulse is a pulse that is followed by a green block indicating a silence of at least Ψ time units. (b) One faulty source that produces spurious pulses. As one of the sources is faulty, two nodes may have different observations on the output of a faulty source given in View 1 (upper figure) and View 2 (lower figure). In this example, a spurious input pulse occurs, that is, an input pulse that is inconsistently detected by different nodes. We devise our construction so that only pulses that follow the frequency bounds of the source are accepted and if the source fails to adhere to the frequency bounds, it will be ignored for some time. Here, the nodes observing View 1 do not receive the input pulse, and thus, they ignore any pulses from the first source for a while (gray region). However, nodes observing View 2 receive a spurious pulse that adheres to the frequency bounds (and do not detect any suspicious behaviour). Thus, nodes with View 2 output a pulse, although nodes observing View 1 do not (red region). Once the faulty source is silenced, the correctly working source has time to produce a good resynchronisation pulse no matter what the faulty source does. This ensures that nodes observing either view generate an output signal within a small enough time window that is followed by at least Ψ time of silence.

output is conservatively filtered out (for a sufficiently large period of time) by a voting mechanism. This happens only for an incorrect output, implying that the fault threshold for the respective instance must have been exceeded. Accordingly, the other instance is operating correctly and, thanks to the absence of further interference from the faulty instance, succeeds in generating a resynchronisation pulse. Figure 10(b) illustrates this idea.

Using Two Pulse Synchronisers. We now overview how to use two pulse synchronisation algorithms to implement our simple resynchronisation algorithm described above. Let

$$n_0 = \lfloor n/2 \rfloor \text{ and } n_1 = \lceil n/2 \rceil$$

$$f_0 = \lfloor (f-1)/2 \rfloor \text{ and } f_1 = \lceil (f-1)/2 \rceil.$$

Observe that we have $n = n_0 + n_1$ and $f = f_0 + f_1 + 1$. We partition the set V of n nodes into two sets V_h for $h \in \{0, 1\}$ such that $V = V_0 \cup V_1$, where $V_0 \cap V_1 = \emptyset$ and $|V_h| = n_h$. We now pick two pulse synchronisation algorithms A_0 and A_1 with the following properties:

- A_h runs on n_h nodes and tolerates f_h faulty nodes,
- A_h stabilises in time $T(A_h)$ and sends $M(A_h)$ bits per time unit and channel, and
- A_h has skew $\sigma \in O(d)$ and accuracy bounds $\Phi_h = (\Phi_h^-, \Phi_h^+)$, where $\Phi_h^-, \Phi_h^+ \in O(\Psi)$.

We let the nodes in set V_h execute the pulse synchronisation algorithm A_h .

An optimistic approach would be to use each A_h as a source of pulses by checking whether at least $n_h - f_h$ nodes in the set V_h generated a pulse within a time window of (roughly) length $\sigma = 2d$. Unfortunately, we cannot directly use the pulse synchronisation algorithms A_0 and A_1 as reliable sources of pulses. There can be a total of $f = f_0 + f_1 + 1 < n/3$ faulty nodes, and thus, it may be that for one $h \in \{0, 1\}$ the set V_h contains more than f_h faulty nodes. Hence, the algorithm A_h may never stabilise and can generate spurious pulses at uncontrolled frequencies. In particular, the algorithm may always generate pulses with frequency less than Ψ , preventing our simple solution from working. However, we are guaranteed that at least one of the algorithms stabilises.

LEMMA 17. *If there are at most f faulty nodes, then there exists $h \in \{0, 1\}$ such that A_h stabilises by time $T(A_h)$.*

PROOF. Observe that $f_0 + f_1 + 1 = f$. In order to prevent both algorithms from stabilising, we need to have at least $f_0 + 1$ faults in the set V_0 and $f_1 + 1$ faults in the set V_1 , totalling $f_0 + f_1 + 2 > f$ faults in the system. $\square$

Once the algorithm A_h for some $h \in \{0, 1\}$ stabilises, we have at least $n_h - f_h$ correct nodes in set V_h locally generate pulses with skew σ and accuracy bounds $\Phi_h = (\Phi_h^-, \Phi_h^+)$. However, it may be that the other algorithm A_{1-h} never stabilises. Moreover, the algorithm A_{1-h} may forever generate spurious pulses at arbitrary frequencies. Here, a spurious pulse refers to any pulse that does not satisfy the skew σ and accuracy bounds of A_{1-h} . For example, a spurious pulse may be a pulse that only a subset of nodes generate, one with too large skew, or a pulse that occurs too soon or too late.

In order to tackle this problem, we employ a series of threshold votes and timeouts to filter out any spurious pulses generated by an unstabilised algorithm *that violate timing constraints*. This way, we can impose some control on the frequency at which an unstabilised algorithm may trigger resynchronisation pulses. As long as these frequency bounds are satisfied, it is inconsequential if a non-stabilised algorithm triggers resynchronisation pulses at a subset of the nodes only. We want our filtering scheme to eventually satisfy the following properties:

- If A_h has stabilised, then all pulses generated by A_h are accepted.
- If A_h has not stabilised, then only pulses that respect given frequency bounds are accepted.

More precisely, in the first case the filtered pulses respect (slightly relaxed) accuracy bounds Φ_h of A_h . In the second case, we enforce that the filtered pulses must either satisfy roughly the same accuracy bounds Φ_h or they must be sufficiently far apart. That is, the nodes will reject any pulses generated by A_h if they occur either too soon or too late.

Once we have the filtering mechanism in place, it becomes relatively easy to implement our conceptual idea for the resynchronisation algorithm. We apply the filtering mechanism for both algorithms A_0 and A_1 and use the filtered outputs as a source of pulses in our algorithm, as illustrated in Figure 11. We are guaranteed that at least one of the sources eventually produces pulses with well-defined accuracy bounds. Furthermore, we also know that the other source must either respect the given accuracy bounds or refrain from generating pulses for a long time. In the case that both sources respect the accuracy bounds, we use the coprimality of frequencies to guarantee a sufficiently large separation window for the resynchronisation pulses. Otherwise,

Fig. 11. Example of the resynchronisation construction for eight nodes tolerating two faults. We partition the network into two parts, each running a pulse synchronisation algorithm A_i . The output of A_i is fed into the respective filter and any pulse that passes the filtering is used as a resynchronisation pulse. The filtering consists of (1) having *all* nodes in the network participate in a threshold vote to see if anyone thinks a pulse from A_i occurred (i.e., enough nodes running A_i generated a pulse) and (2) keeping track of when was the last time a pulse from A_i occurred to check that the accuracy bounds of A_i are respected: pulses that appear too early or too late are ignored. Moreover, if A_i generates pulses at incorrect frequencies, the filtering mechanism blocks all pulses generated by A_i for $\Theta(\Psi)$ time.

Fig. 12. Construction of an f -resilient resynchronisation algorithm on n nodes from f_i -resilient pulse synchronisation algorithms on n_i nodes, where $f = f_0 + f_1 + 1$ and $n = n_0 + n_1$. The n nodes are divided into two groups of n_0 and n_1 nodes. These groups run pulse synchronisation algorithms A_0 and A_1 , respectively. At least one of these algorithms is guaranteed to stabilise eventually. Here, A_1 (gray block) has too many faulty nodes and does not stabilise. All of the n nodes together run two filtering mechanisms F_0 and F_1 for the outputs of A_0 and A_1 , respectively. These ensure that no correct node locally generates a resynchronisation pulse without all correct nodes registering this event, and then apply timeout constraints to enforce the desired frequency bounds.

we exploit the fact that the unreliable source stays silent long enough for the reliable source to generate a pulse with a sufficiently large separation window.

6.2 Filtering Spurious Pulses

Our pulse filtering scheme follows a similar idea as our recent construction of synchronous counting algorithms [24]. However, considerable care is needed to translate the approach from the (much simpler) synchronous round-based model to the bounded-delay model with clock drift. We start by describing the high-level idea of the approach before showing how to implement the filtering scheme in the bounded-delay model considered in this work. Figure 12 illustrates how the underlying pulse synchronisation algorithms are combined with the filtering mechanism.

For convenience, we refer to the nodes in set V_h as *block h* . First, for each block $h \in \{0, 1\}$, every node $v \in G$ performs the following threshold vote:

- (1) If v observes at least $n_h - f_h$ nodes in V_h generating a pulse, vote for generating a resynchronisation pulse.
- (2) If at least $n - f$ nodes in V voted for a pulse by block h (within the time period this should take), then v accepts it.

The idea here is that if some correct node *accepts* a pulse in Step 2, then every correct node must have seen at least $n - 2f \geq f + 1$ votes due to Step (1). Moreover, once a node observes at least $f + 1$ votes, it can deduce that some correct node saw at least $n_h - f_h$ nodes in block h generate a pulse. Thus, if any correct node accepts a pulse generated by block h , then all correct nodes are aware that a pulse *may* have happened.

Second, we have the nodes perform temporal filtering by keeping track of when block h last (may have) generated a pulse. To this end, each node has a local “cooldown timer” that is reset if the node suspects that block h has not yet stabilised. If a pulse is accepted by the above voting mechanism, then a resynchronisation pulse is triggered if the following conditions are met:

- (1) the cooldown timer has expired, and
- (2) not too much time has passed since the most recent pulse from h .

A correct node $v \in G$ resets its cooldown timer if it

- (1) observes at least $f + 1$ votes for a pulse from block h , but not enough time has passed since v last saw at least $f + 1$ votes,
- (2) observes at least $f + 1$ votes, but not $n - f$ votes in a timely fashion, or
- (3) has not observed a pulse from block h for too long, that is, block h should have generated a new pulse by now.

Thus, whenever a block $h \in \{0, 1\}$ triggers a resynchronisation pulse at node $v \in G$, then each node $u \in G$ either resets its cooldown timer or also triggers a resynchronisation pulse. Furthermore, if $v \in G$ does not observe a pulse from block h within the right time window, it will also reset its cooldown counter. Finally, each node refuses to trigger a resynchronisation pulse when its cooldown timer is active. Note that if A_h stabilises, then eventually the cooldown timer for block h expires and is not reset again. This ensures that eventually at least one of the blocks triggers resynchronisation pulses.

Implementation in the Bounded-Delay Model. We implement the threshold voting and temporal filtering with two state machines depicted in Figure 13 and Figure 14. For each block $h \in \{0, 1\}$, every node runs a single copy of the voter and validator state machines in parallel. In the voter state machine given in Figure 13, there are two key states, FAIL and GO, which are used to indicate a local signal for the validator state machine in Figure 14.

The key feature of the voter state machine is the voting scheme: if some node $v \in G$ transitions from VOTE to GO, then all nodes must transition to either FAIL or GO. This is guaranteed by the fact that a node only transitions to GO if it has observed at least $n - f$ nodes in the state VOTE within a short time window. This in turn implies that all nodes must observe at least $n - 2f > f$ nodes in VOTE (in a slightly larger time window). Thus, any node in state IDLE must either transition directly to FAIL or move on to LISTEN. If a node transitions to state LISTEN, then it is bound to either transition to GO or FAIL.

The validator state machine in turn ensures that any subsequent GO transitions of v are at least $T_{\min, h}/\vartheta$ or $T_{\text{cool}}/\vartheta$ time units apart. Moreover, if any FAIL transition occurs at time t , then any subsequent GO transition can occur at time $t + T_{\text{cool}}/\vartheta$ at the earliest. The voter state machine also handles generating a FAIL transition if the underlying pulse synchronisation algorithm does not produce a pulse within time $T_{\max, h}$. This essentially forces the ACT transitions in the validator state

Fig. 13. The voter state machine is the first of the two state machines used to trigger resynchronisation pulses by block $h \in \{0, 1\}$. Every node runs a separate copy of the voter machine for both blocks. The voter state machine performs a threshold vote to ensure that if at *some* node a resynchronisation pulse is (or might be) triggered by block h , then *all* correct nodes see this by observing at least $f + 1$ VOTE messages within T_{att} local time. Note that nodes immediately transition from FAIL and GO to state IDLE, as there are no guards blocking these transitions. The two states are used to signal the validator state machine given in Figure 14 to generate resynchronisation pulses or to refrain from doing so until the cooldown timer expires.

Fig. 14. The validator state machine is the second of the two state machines used to trigger resynchronisation pulses by block $h \in \{0, 1\}$. Every node runs a separate copy of the validator state machine for both blocks. The validator checks that the GO and FAIL transition signals in the voter state machine given in Figure 13 satisfy the minimum time bound and that the GO transitions occur in a timely manner. Note that the transition from state IGNORE to itself resets the timer T_{cool} .

machine to occur between accuracy bounds $\Lambda_h^- \approx T_{min,h}/\vartheta$ and $\Lambda_h^+ \approx T_{max,h}$ or at least T_{cool}/ϑ time apart. Furthermore, if the underlying pulse synchronisation algorithm A_h stabilises, then the ACT transitions roughly follow the accuracy bounds of A_h .

We give a detailed analysis of the behaviour of the two state machines later in Sections 6.5–6.6. The conditions we impose on the timeouts and other parameters are listed in Table 4. As before, the system of inequalities listed in Table 4 can be satisfied by choosing the timeouts carefully. This is done in Section 6.7.

LEMMA 18. *Let σ and $1 < \vartheta < \varphi$ be constants such that $\vartheta^2 \varphi < 31/30$. There exists a constant $\Psi_0(\vartheta, \varphi, d)$ such that, for any given $\Psi > \Psi_0(\vartheta, \varphi, d)$, we can satisfy the constraints in Table 4 by choosing*

Table 4. The Conditions Employed in the Construction of Section 6

(15)	$T_{\min,h} = \Phi_h^- - \rho$
(16)	$T_{\max,h} = \vartheta(\Phi_h^+ + T_{\text{vote}})$
(17)	$T_{\text{vote}} = \vartheta(\sigma + 2d)$
(18)	$T_{\text{idle}} = \vartheta(\sigma + d)$
(19)	$T_{\text{att}} = \vartheta(T_{\text{vote}} + 2d)$
(20)	$\rho = T_{\text{vote}}$
(21)	$T_{\text{cool}} \in \Theta(\max\{\Phi_h^+, \Psi\})$
(22)	$T^* = \max_{h \in \{0,1\}} \{T(\mathbf{A}_h) + 2\Phi_h^+ + T_{\text{cool}} + \sigma + 2d + \rho\}$
(23)	$\Phi_h^- > \Psi + 2\rho$
(24)	$\Phi_h^- \geq T_{\text{vote}} + T_{\text{idle}} + T_{\text{att}} + \sigma + 2d$
(25)	$T_{\min,h} < T_{\text{cool}}$
(26)	$T_{\min,h}/\vartheta > \Psi + \rho$
(27)	$T_{\text{cool}}/\vartheta > 15\beta$
(28)	$C_0 = 4, C_1 = 5$
(29)	$\Lambda_h^+ = T_{\max,h} + T_{\text{vote}}$
(30)	$\Lambda_h^- = T_{\min,h}/\vartheta$
(31)	$\beta > 2\Psi + 4(T_{\text{vote}} + d) + \rho$
(32)	$\beta \cdot (C_h \cdot j) \leq j \cdot \Lambda_h^- < j \cdot \Lambda_h^+ + \rho \leq \beta \cdot (C_h \cdot j + 1)$ for $0 \leq j \leq 3$

Here $h \in \{0, 1\}$.

- (1) $X \in \Theta(\Psi)$,
- (2) $\Phi_0^- = X$ and $\Phi_0^+ = \varphi X$,
- (3) $\Phi_1^- = rX$ and $\Phi_1^+ = \varphi rX$ for a constant $r > 1$, and
- (4) all remaining timeouts in $O(X)$.

The Resynchronisation Algorithm. We now have described all ingredients of our resynchronisation algorithm. It remains to define what is the output of our resynchronisation algorithm. First, for each $h \in \{0, 1\}$, $v \in G$ and $t \geq 0$, let us define an indicator variable for ACT transitions:

$$r_h(u, t) = \begin{cases} 1 & \text{if node } u \text{ transitions to ACT at time } t \\ 0 & \text{otherwise.} \end{cases}$$

Furthermore, for each $v \in V_0 \cup V_1$, we define the output of our resynchronisation algorithm as follows:

$$r(v, t) = \max\{r_0(v, t), r_1(v, t)\}.$$

We say that block h triggers a *resynchronisation pulse* at node v if $r_h(v, t) = 1$. That is, node v generates a resynchronisation pulse if either block 0 or block 1 triggers a resynchronisation pulse at node v . Our goal is to show that there exists a time $t \in O(\Phi^+ + \Psi)$ such that every node $v \in G$ generates a resynchronisation pulse at some time $t' \in [t, t + \rho)$ and neither block triggers a new resynchronisation pulse before time $t + \Psi$ at any node $v \in G$. First, however, we observe that the communication overhead incurred by running the voter and validator state machines is small.

LEMMA 19. *In order to compute the values $r(v, t)$ for each $v \in G$ and $t \geq 0$, the nodes send at most $\max\{M(\mathbf{A}_h)\} + O(1)$ bits per time unit.*

PROOF. For both $h \in \{0, 1\}$, we have every $v \in V_h$ broadcast a single bit when $\mathbf{A}_h$ generates a pulse locally at node v . Observe that we can assume w.l.o.g. that $\mathbf{A}_h$ generates a pulse only once per time unit even during stabilisation: we can design a wrapper for $\mathbf{A}_h$ that filters out any pulses

that occur within e.g., time d of each other. As we have $\Phi_h^- > d$ by Constraint (23), this does not interfere with the pulsing behaviour once the algorithm has stabilised.

In addition to the pulse messages, it is straightforward to verify that the nodes only need to communicate whether they transition to states IDLE, VOTE, or PASS in the voter state machine. Due to the T_{vote} and $T_{\text{max},h}$ timeouts, a node $v \in G$ cannot transition to the same state more than once within d time, as these timeouts are larger than ϑd by Constraint (16) and Constraint (17). $\square$

6.3 Proof of Theorem 3

We take a top-down approach for proving Theorem 3. We delay the detailed analysis of the voter and state machines themselves to later sections and now state the properties we need in order to prove Theorem 3.

First of all, as we are considering self-stabilising algorithms, it takes some time for $\mathbf{A}_h$ to stabilise, and hence, also for the voter and validator state machines to start operating correctly. We will show that this is bound to happen by time

$$T^* \in \max\{T(\mathbf{A}_h)\} + O(\max\{\Phi_h^+\} + T_{\text{cool}}) \subseteq \max\{T(\mathbf{A}_h)\} + O(\Psi).$$

The exact value of T^* is given by Constraint (22) and will emerge later in our proofs. We show in Section 6.4 that the resynchronisation pulses triggered by a single correct block have skew $\rho = T_{\text{vote}} = \vartheta(\sigma + 2d) \in O(d)$ and the desired separation window of length Ψ ; we later argue that a faulty block cannot incessantly interfere with the resynchronisation pulses triggered by the correct block.

LEMMA 20 (STABILISATION OF CORRECT BLOCKS). *Suppose $h \in \{0, 1\}$ is a correct block. Then there exist times $r_{h,0} \in [T^*, T^* + \Phi_h^+ + \rho]$ and for $i \geq 0$ the times $r_{h,i+1} \in [r_{h,i} + \Phi_h^- - \rho, r_{h,i} + \Phi_h^+ + \rho]$ satisfying the following properties for all $v \in G$ and $i \geq 0$:*

- $r_h(v, t) = 1$ for some $t \in [r_{h,i}, r_{h,i} + \rho]$,
- $r_h(v, t') = 0$ for any $t' \in (t, r_{h,i} + \rho + \Psi)$.

We also show that if either block $h \in \{0, 1\}$ triggers a resynchronisation pulse at some correct node $v \in G$, then for every $u \in G$ (1) a resynchronisation pulse is also triggered by h at roughly the same time or (2) node u refrains from generating a resynchronisation pulse for a sufficiently long time. The latter holds true because node u observes that a resynchronisation pulse *might* have been triggered somewhere (due to the threshold voting mechanism); u thus resets its cooldown counter, that is, transitions to state IGNORE in Figure 13. Formally, this is captured by the following lemma, shown in Section 6.5.

LEMMA 21 (GROUPING LEMMA FOR VALIDATOR MACHINE). *Let $h \in \{0, 1\}$ be any block, $t \geq T^*$, and $v \in G$ be such that $r_h(v, t) = 1$. Then there exists $t^* \in [t - 2T_{\text{vote}} - d, t]$ such that, for all $u \in G$, we have*

$$r_{\text{next}}(u, t^*) \in [t^*, t^* + 2(T_{\text{vote}} + d)] \cup [t^* + T_{\text{cool}}/\vartheta, \infty),$$

where $r_{\text{next}}(u, t^*) = \inf\{t' \geq t^* : r_h(u, t') = 1\}$.

Now, with the above two lemmas in mind, we take the following proof strategy. Fix a correct block $k \in \{0, 1\}$ and let $r_{k,0}$ be the time given by Lemma 20. Observe that if no node $v \in G$ has $r_{1-k}(u, t) = 1$ for any $t \in [r_{k,0}, r_{k,0} + \rho + \Psi]$, then all correct nodes succeed in creating a resynchronisation pulse with separation window Ψ . However, it may be the case that the other (possibly faulty) block $1 - k$ spoils the resynchronisation pulse by also triggering a resynchronisation pulse too soon at some correct node. That is, we may have some $v \in G$ that satisfies $r_{1-k}(u, t) = 1$, where $t \in [r_{k,0}, r_{k,0} + \rho + \Psi]$. But then all nodes observe this and the filtering mechanism now

guarantees that the faulty block either obeys the imposed frequency constraints for the following pulses or nodes will ignore them. Either way, we can argue that a correct resynchronisation pulse is generated by the correct block k soon enough.

Accordingly, assume that the faulty block interferes, i.e., generates a spurious resynchronisation pulse at some node $u \in G$ at time $r_{1-k,0} \in (r_{k,0}, r_{k,0} + \rho + \Psi)$. If there is no such time, the resynchronisation pulse of the correct block would have the required separation window of Ψ . Moreover, for all $v \in G$ and $i \geq 0$, we define

$$r_{h,0}(v) = \inf\{t \geq r_{h,0} : r_h(t, v) = 1\}$$

$$r_{h,i+1}(v) = \inf\{t > r_{h,i}(v) : r_h(t, v) = 1\}.$$

Furthermore, for convenience, we define the notation

$$D_h(u) = \begin{cases} \emptyset & \text{if block } h \text{ is correct} \\ [r_{h,0}(u) + T_{\text{cool}}/\vartheta, \infty] & \text{otherwise.} \end{cases}$$

For the purpose of our analysis, we “discretise” our time into chunks of length $\beta \in \Theta(\Psi)$. For any integers $Y > X \geq 0$, we define

$$I_0(X) = r_{k,0} - 2(T_{\text{vote}} + d) + X \cdot \beta,$$

$$I_1(X) = r_{k,0} + 2(T_{\text{vote}} + d) + \Psi + X \cdot \beta,$$

$$I(X, Y) = [I_0(X), I_1(Y)).$$

We abbreviate $\Lambda_h^- = T_{\min,h}/\vartheta$ and $\Lambda_h^+ = T_{\max,h} + T_{\text{vote}}$ (Constraint (29) and Constraint (30)). The following lemma is useful for proving Theorem 3. We defer its proof to Section 6.6.

LEMMA 22 (RESYNCHRONISATION FREQUENCY). *For any $i \geq 0$ and $v \in G$, it holds that*

$$r_{h,i}(v) \in [r_{h,0} - 2(T_{\text{vote}} + d) + i \cdot \Lambda_h^-, r_{h,0} + 2(T_{\text{vote}} + d) + i \cdot \Lambda_h^+) \cup D_h(v).$$

COROLLARY 8. *Let $h \in \{0, 1\}$ and $0 \leq i \leq 3$. For any $v \in G$, we have*

$$r_{h,i}(v) \in I(i \cdot C_h, i \cdot C_h + 1) \cup D_h(v).$$

PROOF. Recall that $r_{1-k,0} \in (r_{k,0} - 2(T_{\text{vote}} + d), r_{k,0} + \rho + \Psi)$. By Constraint (32), for all $i \leq 3$ and $h \in \{0, 1\}$, it holds that

$$\beta \cdot C_h \cdot i \leq i \cdot \Lambda_h^- < i \cdot \Lambda_h^+ + \rho \leq \beta \cdot (C_h \cdot i + 1).$$

Using Lemma 22, this inequality, and the above definitions, a straightforward manipulation shows that $r_{h,j}(v)$ lies in the interval

$$\begin{aligned} & [r_{h,0} - 2(T_{\text{vote}} + d) + i \cdot \Lambda_h^-, r_{h,0} + 2(T_{\text{vote}} + d) + i \cdot \Lambda_h^+] \cup D_h(v) \\ & \subseteq [r_{k,0} - 2(T_{\text{vote}} + d) + i \cdot \Lambda_h^-, r_{k,0} + 2(T_{\text{vote}} + d) + \Psi + \rho + i \cdot \Lambda_h^+] \cup D_h(v) \\ & \subseteq [r_{k,0} - 2(T_{\text{vote}} + d) + \beta \cdot (C_h \cdot i), r_{k,0} + 2(T_{\text{vote}} + d) + \Psi + \beta \cdot (C_h \cdot i + 1)] \cup D_h(v) \\ & = [I_0(i \cdot C_h), I_1(i \cdot C_h + 1)) \cup D_h(v) \\ & = I(i \cdot C_h, i \cdot C_h + 1) \cup D_h(v). \end{aligned}$$

□

With the above results, we can now show that eventually the algorithm outputs a good resynchronisation pulse.

LEMMA 23. *There exists a time $t \in \max\{A_h\} + O(\Psi)$ such that for all $v \in G$ there exists a time $t_v \in [t, t + \rho]$ satisfying*

- $r(v, t_v) = 1$, and
- $r(v, t') = 0$ for $h \in \{0, 1\}$ and any $t' \in (t_v, t_v + \Psi)$.

PROOF. Suppose block $k \in \{0, 1\}$ is correct. The lemma follows by proving that we have the following properties for some $t \leq I_1(11)$ and each $v \in G$:

- $r_k(v, t_v) = 1$ for some $t_v \in [t, t + \rho]$, and
- $r_h(v, t') = 0$ for $h \in \{0, 1\}$ and any $t' \in (t_v, t_v + \Psi)$.

Recall that $r_{1-k,0} \in (r_{k,0} - 2(T_{\text{vote}} + d), r_{k,0} + \rho + \Psi)$, as otherwise the claim trivially follows for $t = r_{k,0}$. Consider any $v \in G$. Corollary 8 and the fact that $C_0 = 4$ by Constraint (28) imply

$$\begin{aligned} r_{0,0}(v) &\in I(0, 1) \cup D_0(v), \\ r_{0,1}(v) &\in I(4, 5) \cup D_0(v), \\ r_{0,2}(v) &\in I(8, 9) \cup D_0(v), \\ r_{0,3}(v) &\in I(12, 13) \cup D_0(v). \end{aligned}$$

As $T_{\text{cool}}/\vartheta \geq 15\beta$ by Constraint (27), it follows for all $t \geq r_{0,0}$ that if $r_0(v, t) = 1$, then

$$t \in I(0, 1) \cup I(4, 5) \cup I(8, 9) \cup [I_0(12), \infty].$$

Similarly, as $C_1 = 5$, for all $t \geq r_{1,0}$ we get that if $r_1(v, t) = 1$, then

$$t \in I(0, 1) \cup I(5, 6) \cup I(10, 11) \cup [I_0(15), \infty].$$

Let k be the correct block we have fixed. Recall that $D_k(v) = \emptyset$. The claim now follows from a simple case analysis:

- (1) If $k = 0$, then $r_{k,2}(v) \in I(8, 9)$ and $r_{1-k}(v, t') = 0$ for all $t' \in [I_1(6), I_0(10)) \supset [I_0(8), I_1(9) + \Psi + \rho]$ (by Constraint (31)).
- (2) If $k = 1$, then $r_{k,2}(v) \in I(10, 11)$ and $r_{1-k}(v, t') = 0$ for all $t' \in [I_1(9), I_0(12)) \supset [I_0(10), I_1(11) + \Psi + \rho]$ (by Constraint (31)).

Thus, in both cases, $t = \min_{v \in G} \{r_{k,2}(v)\}$ satisfies the claim of the lemma, provided that $t \leq I_1(11) \in \max\{\mathbf{A}_h\} + O(\Psi)$. This is readily verified from the constraints given in Table 4. $\square$

THEOREM 3. Let $f, n_0, n_1 \in \mathbb{N}$ and $1 < \vartheta \leq 1.004$. Define

$$n = n_0 + n_1, \quad f_0 = \lfloor (f - 1)/2 \rfloor, \quad f_1 = \lceil (f - 1)/2 \rceil.$$

For any $\Psi \in \Omega(1)$ and sufficiently small constant $\varphi > \varphi_0(\vartheta)$, there exists a bound $T_0 \in \Theta(\Psi)$ such that the following claim holds. If, for both $i \in \{0, 1\}$, there exists pulse synchronisation algorithm $\mathbf{A}_i$ that

- runs on n_i nodes and has resilience f_i ,
- has skew $\sigma = 2d$, and
- has accuracy bounds $\Phi_i^- = T$ and $\Phi_i^+ = T\varphi$, where $T_0 \leq T$ and $T \in O(\Psi)$,

then there exists a resynchronisation algorithm $\mathbf{B}$ that

- runs on n nodes and has resilience f ,
- has skew $\rho \in O(d)$ and separation window of length Ψ ,
- generates a resynchronisation pulse by time $T(\mathbf{B}) \in \max\{T(\mathbf{A}_0), T(\mathbf{A}_1)\} + O(\Psi)$, and
- has nodes send $M(\mathbf{B}) \in \max\{M(\mathbf{A}_0), M(\mathbf{A}_1)\} + O(1)$ bits per time unit and channel.

PROOF. Computation shows that for $\vartheta \leq 1.004$, we have that $\vartheta^2 \varphi_0(\vartheta) < 31/30$, where $\varphi_0(\vartheta) = 1 + 5(\vartheta - 1)/(2 + 2\vartheta - 3\vartheta^2)$ is given in Corollary (7). Thus, Lemma 18 shows that for a sufficiently small choice of $\varphi > \varphi_0(\vartheta) > \vartheta$, we can pick $\Phi_h^- \in \Theta(\Psi)$, where $1 < \max\{\Phi_h^+/\Phi_h^-\} \leq \varphi$, such that the conditions given in Table 4 are satisfied. Thus, using our assumption, we can choose the algorithms $\mathbf{A}_h$ with these accuracy bounds Φ_h ; note that here we use sufficiently small ϑ and φ that satisfy both our initial assumption and the preconditions of Lemma 18.

In order to compute the output value $r(v, t) \in \{0, 1\}$ for each $v \in G$ and $t \geq 0$, Lemma 19 shows that our resynchronisation algorithm only needs to communicate $O(1)$ bits per time unit in addition to the message sent by underlying pulse synchronisation algorithms A_0 and A_1 . By Lemma 23, we have that a good resynchronisation pulse with skew $\rho \in O(d)$ happens at a time $t \in \max\{A_h\} + O(\Psi)$. $\square$

6.4 Proof of Lemma 20

We now show that eventually a correct block $h \in \{0, 1\}$ will start triggering resynchronisation pulses with accuracy bounds $\Lambda_h = (\Lambda_h^-, \Lambda_h^+)$. Our first goal is to show that after the algorithm A_h has stabilised in a correct block h , all correct nodes will start transitioning to GO in a synchronised fashion. Then we argue that eventually the transitions to the state GO will be coupled with the transitions to ACT.

Recall that the pulse synchronisation algorithm A_h has skew σ and accuracy bounds $\Phi_h = (\Phi_h^-, \Phi_h^+)$. Let $p_h(v, t) \in \{0, 1\}$ indicate whether node $v \in V_h \setminus F$ generates a pulse according to algorithm A_h at time t . If block $h \in \{0, 1\}$ is correct, then by time $T(A_h)$ the algorithm A_h has stabilised. Moreover, then there exists a time $T(A_h) \leq p_{h,0} \leq T(A_h) + \Phi_h^+$ such that each $v \in V_h \setminus F$ satisfies $p_h(v, t) = 1$ for some $t \in [p_{h,0}, p_{h,0} + \sigma)$. Since block h and algorithm A_h are correct, there exist for each $v \in V_h \setminus F$ and $i \geq 0$ the following values:

$$\begin{aligned} p_{h,i}(v) &= \inf\{t \geq p_{h,i} : p_h(v, t) = 1\} \neq \infty, \\ p_{h,i+1} &\in [p_{h,i} + \Phi_h^-, p_{h,i} + \Phi_h^+), \\ p_{h,i+1}(v) &\in [p_{h,i+1}, p_{h,i+1} + \sigma). \end{aligned}$$

That is, A_h generates a pulse at node $v \in V_h$ for the i th time after stabilisation at time $p_{h,i}(v)$.

First, let us observe that the initial “clean” pulse makes every correct node transition to GO or FAIL, thereby resetting the $T_{\max,h}$ timeouts, where the nodes will wait until the next pulse.

LEMMA 24. *Suppose block $h \in \{0, 1\}$ is correct. Each correct node $v \in G$ is in state IDLE at time $p_{h,1}$ and its local $T_{\max,h}$ timer does not expire during the interval $[p_{h,1}, p_{h,1} + \sigma + d)$.*

PROOF. First, observe that if the timer $T_{\max,h}$ is reset at time $p_{h,0}$ or later, then it will not expire before time $p_{h,0} + T_{\max,h}/\vartheta > p_{h,0} + \Phi_h^+ + \sigma + d \geq p_{h,1} + \sigma + d$ by Constraint (16) and Constraint (17). Because every node receives $n_h - f_h$ pulse messages from different nodes in V_h during $(p_{h,0}, p_{h,0} + \sigma + d)$ and $T_{\text{idle}} = \vartheta(\sigma + d)$ by Constraint (18), every node that is in state IDLE at time $p_{h,0}$ leaves this state during $(p_{h,0}, p_{h,0} + \sigma + d)$. Recall that nodes cannot stay in states GO or FAIL. Because nodes leave states LISTEN, VOTE, and PASS when the timeout $\langle T_{\text{vote}}, \text{LISTEN} \rangle$ expires, any node not in state IDLE must transition to this state within T_{vote} time. Accordingly, each correct node resets its $T_{\max,h}$ timer during $(p_{h,0}, p_{h,0} + T_{\text{vote}} + \sigma + d)$.

Next, note that during $(p_{h,0} + T_{\text{idle}} + \sigma + d, p_{h,1})$, no correct node has any pulse messages from correct nodes in V_h in its respective buffer. Therefore, no correct node can transition to vote during this time interval. This, in turn, implies that no correct node has any VOTE messages from correct nodes in its respective buffer with timeout T_{att} during $(p_{h,0} + T_{\text{idle}} + T_{\text{att}} + \sigma + 2d, p_{h,1})$. Therefore, no correct node can leave state IDLE during this period. Finally, any correct nodes not in state IDLE at time $p_{h,0} + T_{\text{idle}} + T_{\text{att}} + \sigma + 2d$ must transition back to IDLE by time $p_{h,0} + T_{\text{vote}} + T_{\text{idle}} + T_{\text{att}} + \sigma + 2d$. As $T_{\text{vote}} + T_{\text{idle}} + T_{\text{att}} + \sigma + 2d \leq \Phi_h^-$ by Constraint (24), the claim follows. $\square$

Let us define an indicator variable for GO transitions:

$$g_h(v, t) = \begin{cases} 1 & \text{if node } u \text{ transitions to GO at time } t \\ 0 & \text{otherwise.} \end{cases}$$

Similarly to above, we now also define for $v \in G$ and $i \geq 1$ the values

$$\begin{aligned} g_{h,1} &= \inf\{t \geq p_1 : g_h(u, t) = 1, u \in G\}, \\ g_{h,i}(v) &= \inf\{t \geq g_{h,1} : g_h(v, t) = 1\}, \\ g_{h,i+1}(v) &= \inf\{t > g_{h,i}(v)\}. \end{aligned}$$

In words, the time $g_{h,1}$ is the minimal time that some correct node transitions to state GO in the voter state machine of block h at or after the second pulse of A_h after stabilisation. The two other values indicate the i th time a correct node $v \in G$ transitions to GO starting from time $g_{h,1}$.

We now show that starting from the second pulse $p_{h,1}$ of a correct block h , the GO signals essentially just “echo” the pulse.

LEMMA 25. *If block h is correct, then, for all $v \in G$ and $i > 0$, it holds that*

$$g_{h,i}(v) \in (p_{h,i} + \sigma + 2d, p_{h,i} + \sigma + 2d + \rho),$$

where $\rho = \vartheta(\sigma + 2d) = T_{\text{vote}}$. Moreover, node v does not transition to state FAIL at any time $t \geq p_{h,1}$.

PROOF. By Lemma 24, we have that each $v \in G$ is in state IDLE at time $p_{h,1}$ and Guard G3 is not active during $[p_{h,1}, p_{h,1} + \sigma + d)$. During $(p_{h,1}, p_{h,1} + \sigma + d)$, i.e., within T_{idle} local time by Constraint (18), each node receives $n_h - f_h$ pulse messages from different nodes in V_h and thus transitions to VOTE. Thus, all nodes receive $n - f$ VOTE messages from different nodes during $(p_{h,1}, p_{h,1} + \sigma + 2d)$. As $T_{\text{vote}}/\vartheta = \sigma + 2d$ by Constraint (17), each correct node transitions to PASS before $\langle T_{\text{vote}}, \text{LISTEN} \rangle$ expires and transitions to GO at a time from $(p_{h,1} + \sigma + 2d, p_{h,1} + \sigma + 2d + T_{\text{vote}}) = (p_{h,1} + \sigma + 2d, p_{h,1} + \sigma + 2d + \rho)$. In particular, it resets its buffers after all pulse and VOTE messages from correct nodes are received. Consequently, correct nodes stay in state IDLE until either the next pulse or $T_{\text{max},h}$ expires. The former occurs at all correct nodes in V_h no earlier than time $p_{h,2} > p_{h,1} + T_{\text{att}} + \sigma + 2d > p_{h,1} + T_{\text{vote}} + \sigma + 2d$ by Constraint (19) and Constraint (24), and each pulse is received before time $p_{h,2} + \sigma + d < p_{h,1} + \Phi_h^+ + T_{\text{vote}}$ by Constraint (16) and Constraint (17). Thus, each correct node stay in state IDLE until time $p_{h,2}$ with $T_{\text{max},h}$ expiring no earlier than time $p_{h,2} + \sigma + d$. Consequently, we can repeat the above reasoning inductively to complete the proof. $\square$

We define the following time bound for each $h \in \{0, 1\}$:

$$T_h^* = T(A_h) + T_{\text{cool}} + 2\Phi_h^+ + \sigma + 2d + \rho.$$

We now show that by time T_h^* we are guaranteed that transitions to GO and ACT have become coupled if block h is correct.

LEMMA 26. *Suppose block $h \in \{0, 1\}$ is correct. Then, for any $v \in G$ and $t \geq T_h^*$, it holds that*

$$r_h(v, t) = 1 \text{ if and only if } g_h(v, t) = 1.$$

PROOF. Note that node $u \in G$ can transition to state ACT from WAIT in the validator state machine only if the voter state machine transitions to GO. Consider the time $g_{h,i}(v)$ for $i > 0$. Observe that there are three states in which v may be at this time: WAIT, HOLD, or IGNORE. We argue that, in each case, node v eventually is in state WAIT in the validator state machine when the voter state machine transitions to state GO, and thus, node v transitions to ACT.

First of all, note that, by Lemma 25, node v does not transition to the state FAIL at any time $t \geq p_{h,1}$. We utilise this fact in each of the three cases below:

- (1) In the first case, node v transitions from WAIT to ACT at time $g_{h,i}(v)$ and hence we have $r_h(v, g_{h,i}(v)) = 1$. By applying both Lemma 25 and Constraint (15), we get that $g_{h,i+1}(v) \geq g_{h,i}(v) + \Phi_h^- \geq g_{h,i}(v) + T_{\text{min},h}$. Moreover, v does not transition to the FAIL state in the

voter state machine at any time $t \geq g_{h,j}$. Hence, by induction, node v transitions from state WAIT to ACT at time $g_{h,j}(v)$ for each $j \geq i$.

- (2) In the second case, v transitions from HOLD to IGNORE at time $g_{h,i}(v)$. By time $r \leq g_{h,i}(v) + T_{\text{cool}}$, node v transitions to WAIT. Hence, for any j with $g_{h,j}(v) \geq r$, the first case applies.
- (3) In the third case, v resets its T_{cool} timeout and remains in state IGNORE until, at a time $r \leq g_{h,i}(v) + T_{\text{cool}}$, its T_{cool} timer expires and v transitions to WAIT. Again, for any j with $g_{h,j}(v) \geq r$, the first case applies.

Now consider the time $g_{h,1}(v) \in [p_{h,1} + \sigma + 2d, p_{h,1} + \sigma + 2d + \rho]$ given by Lemma 25. From the above case analysis, we get that node v is in state WAIT by time

$$g_{h,1}(v) + T_{\text{cool}} < p_{h,1} + \sigma + 2d + \rho + T_{\text{cool}} \leq T_h^*,$$

and from then on each transition to GO entails a transition to ACT, as claimed. $\square$

LEMMA 20 (STABILISATION OF CORRECT BLOCKS). *Suppose $h \in \{0, 1\}$ is a correct block. Then there exist times $r_{h,0} \in [T^*, T^* + \Phi_h^+ + \rho]$ and for $i \geq 0$ the times $r_{h,i+1} \in [r_{h,i} + \Phi^- - \rho, r_{h,i} + \Phi^+ + \rho]$ satisfying the following properties for all $v \in G$ and $i \geq 0$:*

- $r_h(v, t) = 1$ for some $t \in [r_{h,i}, r_{h,i} + \rho]$,
- $r_h(v, t') = 0$ for any $t' \in (t, r_{h,i} + \rho + \Psi)$.

PROOF. First, observe that, by Constraint (22), we have $T^* = \max\{T_h^*\}$. Let $p_{h,j} \in [T^* - \sigma - 2d, T^* - \sigma - 2d + \Phi_h^+]$ for some $j > 0$. By Lemma 25, we have that for all $i \geq j$ and $v \in G$ it holds that

$$g_{h,i}(v) \in (p_{h,i} + \sigma + 2d, p_{h,i} + \sigma + 2d + \rho),$$

and by Lemma 26, we have $r_h(v, t) = 1$ for all $g_{h,i}(v) = t \geq T^*$ and $r_h(v, t) = 0$ for all other $t \geq T^*$. We set $r_{h,i} = \min_{v \in G} \{g_{h,j+i}(v)\}$. As $p_{h,i'+1} - p_{h,i'} \in [\Phi^-, \Phi^+]$ for all $i' \geq 0$, this shows all required time bounds but $r_h(v, t') = 0$ for each $v \in G$, i , and $t' \in (g_{h,j+i}(v), r_{h,i} + \rho + \Psi)$. The latter follows because $\Phi^- > \Psi + 2\rho$ by Constraint (23). $\square$

6.5 Proof of Lemma 21

LEMMA 21 (GROUPING LEMMA FOR VALIDATOR MACHINE). *Let $h \in \{0, 1\}$ be any block, $t \geq T^*$, and $v \in G$ be such that $r_h(v, t) = 1$. Then there exists $t^* \in [t - 2T_{\text{vote}} - d, t]$ such that, for all $u \in G$, we have*

$$r_{\text{next}}(u, t^*) \in [t^*, t^* + 2(T_{\text{vote}} + d)] \cup [t^* + T_{\text{cool}}/\vartheta, \infty),$$

where $r_{\text{next}}(u, t^*) = \inf\{t' \geq t^* : r_h(u, t') = 1\}$.

In order to show Lemma 21, we analyse how the voting and validator state machines given in Figure 13 and Figure 14 behave. We show that the voter machines for a single block $h \in \{0, 1\}$ are roughly synchronised in the following sense: if some correct node transitions to GO, then every correct node will transition to either GO or FAIL within a short time window.

LEMMA 27. *Let $t \geq 2T_{\text{vote}} + d$ and $v \in G$ such that $g_h(v, t) = 1$. Then there exists $t^* \in (t - 2T_{\text{vote}} - d, t]$ such that all correct nodes transition to GO or FAIL during the interval $[t^*, t^* + 2(T_{\text{vote}} + d)]$.*

PROOF. Note that, at any time $t' \geq T_{\text{vote}} + d$, any VOTE message stored at a correct node (supposedly) sent by another correct node must actually have been sent at a time greater than 0. Since $v \in G$ satisfies $g_h(v, t) = 1$, this means that it transitioned to LISTEN at a time $t' \in [t - T_{\text{vote}}, t]$, implying that v received at least $n - 2f$ VOTE messages from different correct nodes during the interval $[t' - T_{\text{vote}}, t']$. Hence, every correct $u \in G$ must receive at least $n - 2f > f$ VOTE messages from different nodes during the interval $I = (t' - T_{\text{vote}} - d, t' + d)$.

Let $t^* < t'$ be the minimal time a correct node transitions to `VOTE` during the interval I . Consider any node $u \in G$. By the above observations, u has stored at least $f + 1$ `VOTE` messages in the buffer using timeout T_{att} at some time $t'' \in [t^*, t' + d]$ (where we use Constraint (19)) and must transition to `LISTEN` in case it is in state `IDLE`. Any node that is not in state `IDLE` will transition to `GO` or `FAIL` within T_{vote} time by Guard G4. Overall, each correct node must transition to `FAIL` or `GO` during the interval $[t^*, t' + T_{\text{vote}} + d] \subseteq [t^*, t^* + 2(T_{\text{vote}} + d)]$. $\square$

We now show a similar synchronisation lemma for the validator state machines as well: if some correct node transitions to `ACT` and triggers a resynchronisation pulse, then every correct node triggers a resynchronisation pulse or transitions to `IGNORE` within a short time window.

LEMMA 28. *Let $t \geq 2T_{\text{vote}} + d$ and suppose $r_h(u, t) = 1$ for some $v \in G$. Then there exists a time $t^* \in (t - 2T_{\text{vote}} - d, t]$ such that all correct nodes transition to `ACT` or `IGNORE` during the time interval $[t^*, t^* + 2(T_{\text{vote}} + d)]$.*

PROOF. Suppose some node v transitions to state `ACT` at time t . Then it must have transitioned to state `GO` in the voter state machine at time t as well. By Lemma 27, we get that there exists $t^* \in (t - 2T_{\text{vote}} - d, t]$ such that all correct nodes transition to `GO` or `FAIL` during the interval $[t^*, t^* + 2(T_{\text{vote}} + d)]$. Once $u \in G$ transitions to either of these states in the voter state machine, this causes a transition in the validator state machine to either `ACT` or `IGNORE`, as can be seen from Figure 14. Hence, during the same interval, all correct nodes will either transition to `ACT` or `IGNORE`. $\square$

Observe that once $v \in G$ transitions to `IGNORE` at time t , then it cannot transition back to `ACT` before T_{cool} time has passed on its local clock, that is, before time $t + T_{\text{cool}}/\vartheta$. Thus, Lemma 28 now implies Lemma 21, as $T^* \geq 2T_{\text{vote}} + d$.

6.6 Proof of Lemma 22

We now aim to prove Lemma 22. Hence, let $r_{h,0}$ be as defined in Section 6.3. We have shown above that if block h is correct, then the resynchronisation pulses generated by block h are coupled with the pulses generated by the underlying pulse synchronisation algorithm A_h . We will argue that any block $h \in \{0, 1\}$, including a faulty one, must either respect the accuracy bounds $\Lambda_h = (\Lambda_h^-, \Lambda_h^+)$ when triggering resynchronisation pulses or refrain from triggering a resynchronisation pulse for at least time $T_{\text{cool}}/\vartheta$.

LEMMA 29. *Let $u \in G$, $h \in \{0, 1\}$, and $i \geq 0$. Then*

$$r_{h,i+1}(u) \in [r_{h,i}(u) + \Lambda_h^-, r_{h,i}(u) + \Lambda_h^+] \cup D_h(u).$$

PROOF. First, observe that in case $r_{h,i}(u) = \infty$ for any $i \geq 0$, by definition $r_{h,i+1}(u) = \infty \in D_h(u)$ and the claim holds.

Hence, let $t = r_{h,i}(u) \neq \infty$. Since u transitions to `ACT` at time t , it follows that u also transitions to state `GO` at time t , that is, $g_h(u, t) = 1$. Therefore, u will transition to state `IDLE` in the voter state machine and state `HOLD` in the validator state machine. Observe that u cannot transition to `ACT` again before time $t + \min\{T_{\text{min},h}, T_{\text{cool}}\}/\vartheta$, that is, before either local timer $T_{\text{min},h}$ or T_{cool} expires. Since $T_{\text{min},h} < T_{\text{cool}}$ by Constraint (25), we get that $r_{h,i+1}(u) \geq t + T_{\text{min},h}/\vartheta = t + \Lambda_h^-$.

Next note that u transitions to `FAIL` when (1) the local timer $T_{\text{max},h}$ expires when in state `IDLE` or (2) T_{vote} , which is reset only upon leaving `IDLE`, expires when not in state `IDLE`. Thus, by time $t + T_{\text{max},h} + T_{\text{listen}}$, node u transitioned to `FAIL` or `GO` again. This implies that by time $t + T_{\text{max},h} + T_{\text{vote}}$ node u has transitioned to either `ACT` or `IGNORE` in the validator state machine. Hence, we get that $r_{h,i+1}(u) \leq t + T_{\text{max},h} + T_{\text{vote}} = r_{h,i}(u) + \Lambda_h^+$ or $r_{h,i+1}(u) \geq r_{h,i}(u) + T_{\text{cool}}/\vartheta \geq r_{h,0}(u) + T_{\text{cool}}/\vartheta$. Therefore, the claim follows. $\square$

Lemma 22 readily follows.

LEMMA 22 (RESYNCHRONISATION FREQUENCY). *For any $i \geq 0$ and $v \in G$, it holds that*

$$r_{h,i}(v) \in [r_{h,0} - 2(T_{\text{vote}} + d) + i \cdot \Lambda_h^-, r_{h,0} + 2(T_{\text{vote}} + d) + i \cdot \Lambda_h^+] \cup D_h(v).$$

PROOF. Again, observe that if $r_{h,i}(v) = \infty$, then the claim vacuously holds for all $j \geq i$. We prove the claim by induction on increasing i , so w.l.o.g. we may assume that $r_{h,i} \neq \infty$ for all $i \geq 0$. The base case $i = 0$ follows directly from the definition of $r_{h,0}$ and Lemma 21. By applying Lemma 29 to indexes i , v , and h , and using the induction hypothesis, we get that $r_{h,i+1}(v)$ lies in the interval

$$\begin{aligned} & [r_{h,i}(v) + \Lambda_h^-, r_{h,i}(v) + \Lambda_h^+] \cup [r_{h,i}(v) + T_{\text{cool}}/\vartheta, \infty] \\ & \subseteq [r_{h,0} - 2(T_{\text{vote}} + d) + (i+1) \cdot \Lambda_h^-, r_{h,0} + 2(T_{\text{vote}} + d) + (i+1) \cdot \Lambda_h^+] \cup D_h(v). \quad \square \end{aligned}$$

6.7 Proof of Lemma 18

LEMMA 18. *Let σ and $1 < \vartheta < \varphi$ be constants such that $\vartheta^2\varphi < 31/30$. There exists a constant $\Psi_0(\vartheta, \varphi, d)$ such that, for any given $\Psi > \Psi_0(\vartheta, \varphi, d)$, we can satisfy the constraints in Table 4 by choosing*

- (1) $X \in \Theta(\Psi)$,
- (2) $\Phi_0^- = X$ and $\Phi_0^+ = \varphi X$,
- (3) $\Phi_1^- = rX$ and $\Phi_1^+ = \varphi rX$ for a constant $r > 1$, and
- (4) all remaining timeouts in $O(X)$.

PROOF. We show that we can satisfy the constraints by setting

$$\begin{aligned} \Phi_0^- &= X \\ \Phi_0^+ &= \varphi X \\ \Phi_1^- &= rX \\ \Phi_1^+ &= r\varphi X \\ \Psi &= aX \\ \beta &= bX \\ T_{\text{cool}} &= cX, \end{aligned}$$

where $a = b/3$, $b = 6/25 \cdot \vartheta\varphi$, $c = 16\vartheta b$, and $r = 31/25$, and by picking a sufficiently large $X > X_0(\vartheta, \varphi, d)$. Here $X_0(\vartheta, \varphi, d)$ depends only on the given constants. Note that the choice of T_{cool} satisfies Constraint (21).

First, let us pick the values for the remaining timeouts and variables as given by Constraints (15)–(20), (22), and (28)–(30); it is easy to check that these equalities can be satisfied simultaneously. Regarding Constraint (23), observe that $2/25 \cdot \vartheta\varphi < 1$ and

$$\Phi_h^- \geq X > 2/25 \cdot \vartheta\varphi X + 2\rho = aX + 2\rho = \Psi + 2\rho$$

when $X > 2\rho/(1 - 2/25 \cdot \vartheta\varphi) = 2\vartheta(\sigma + 2d)/(1 - 2/25 \cdot \vartheta\varphi)$, that is, X is larger than a constant. Furthermore, Constraint (24) is also satisfied by picking the constant bounding X from below to be large enough.

To see that Constraint (25) holds, observe that $T_{\text{cool}} = cX = 16\vartheta^2\varphi X \cdot 6/25 > 96/25 \cdot X > 31/25 \cdot X = rX \geq \Phi_h^-$ for both $h \in \{0, 1\}$. Assuming $X > 2\rho \cdot 375/344 = 2\rho/(1 - 2/25 \cdot 31/30)$, Constraint (26) is satisfied since

$$\Phi_h^- - \rho \geq X - \rho > 2/25 \cdot 31/30 \cdot X + \rho > 2/25 \cdot \vartheta^2\varphi X + \rho > b/3 \cdot X + \rho = aX + \rho = \Psi + \rho.$$

Constraint (27) is satisfied as $T_{\text{cool}}/\vartheta = cX/\vartheta = 16bX = 16\beta > 15\beta$. Having $X > 3/b \cdot (5\rho + 4d)$ yields that Constraint (31) is satisfied, since then

$$2\Psi + 4(T_{\text{vote}} + d) + \rho = 2\Psi + 5\rho + 4d = 2b/3 \cdot X + 5\rho + 4d < bX = \beta.$$

It remains to address Constraint (32). As Constraint (15) and Constraint (30) hold, the first inequality of Constraint (32) is equivalent to

$$\vartheta\beta C_h = 6/25 \cdot \vartheta^2\varphi X C_h \leq \Phi_h^- - \rho.$$

We have set $C_0 = 4$ in accordance with Constraint (28). For $X > \rho/(1 - 24/25 \cdot \vartheta^2\varphi)$,

$$24/25 \cdot \vartheta^2\varphi X < X - \rho = \Phi_0^- - \rho$$

thus shows that the inequality holds. Concerning $h = 1$, we set $C_1 = 5$. Recalling that $\vartheta^2\varphi < 31/30$, we may assume that $X > \rho/(31/25 - 6/5 \cdot \vartheta^2\varphi)$, yielding

$$6/5 \cdot \vartheta^2\varphi X < 31/25 \cdot X - \rho = rX - \rho = \Phi_1^- - \rho,$$

i.e., the first inequality of Constraint (32) is satisfied for $h = 1$. The middle inequality is trivially satisfied, as $\Lambda_h^- < \Lambda_h^+$. By the already established equalities, the final inequality in Constraint (32) is equivalent to

$$j\vartheta\Phi_h^+ + ((\vartheta + 1)j + 1)\rho \leq \beta(C_h \cdot j + 1)$$

for all $h \in \{0, 1\}$ and $0 \leq j \leq 3$.

Let $A_j = ((\vartheta + 1)j + 1)\rho$ and observe that $25/3 \cdot A_3 > 25/4 \cdot A_2 > 5A_1$. For any $X > 25/3 \cdot A_3$ and $h = 0$, a simple calculation thus shows

$$\vartheta\varphi X + A_1 < 30/25 \cdot \vartheta\varphi X = 5bX,$$

$$2\vartheta\varphi X + A_2 < 54/25 \cdot \vartheta\varphi X = 9bX,$$

$$3\vartheta\varphi X + A_3 < 78/25 \cdot \vartheta\varphi X = 13bX.$$

Since $\Phi_0^+ = \varphi X$, $\beta = bX$, and $C_0 = 4$, this covers the case of $h = 0$. Similarly, as $r = 31/25 = 1 + b/(\vartheta\varphi)$, we have

$$\vartheta\varphi rX + A_1 < 6bX,$$

$$2\vartheta\varphi rX + A_2 < 11bX,$$

$$3\vartheta\varphi rX + A_3 < 16bX,$$

covering the case of $h = 1$ with $C_1 = 5$. Overall, we conclude that Constraint (32) is satisfied.

Finally, observe that, in all cases, we assumed that X is bounded from below by a function $X_0(\vartheta, \varphi, d)$ that depends only on the constants ϑ , φ , and d . Thus, the constraints can be satisfied by picking $X > X_0(\vartheta, \varphi, d)$ which yields that we can satisfy the constraints for any $\Psi > \Psi_0(\vartheta, \varphi, d) = aX_0(\vartheta, \varphi, d)$. $\square$

7 RANDOMISED ALGORITHMS

While we have so far only considered deterministic algorithms, our framework also extends to randomised algorithms. In particular, this allows us to obtain faster algorithms by simply replacing the synchronous consensus algorithms we use by randomised variants. Randomised consensus algorithms can break the linear-time lower bound [19] for deterministic algorithms [21, 29]. This in turn allows us to construct the first pulse synchronisation algorithms that stabilise in sublinear time.

Typically, when considering randomised consensus, one relaxes the termination property: it suffices that the algorithm terminates with probability 1 and gives probabilistic bounds on the (expected, w.h.p., etc.) round complexity. However, our framework operates based on a deterministic termination guarantee, where the algorithm is assumed to declare its output in R rounds.

Therefore, we instead relax the *agreement* property so that it holds with a certain probability only. Formally, node v is given an input $x(v) \in \{0, 1\}$, and it must output $y(v) \in \{0, 1\}$ such that the following properties hold:

- (1) *Agreement.* With probability at least p , there exists $y \in \{0, 1\}$ such that $y(v) = y$ for all correct nodes v .
- (2) *Validity.* If, for $x \in \{0, 1\}$, it holds that $x(v) = x$ for all correct nodes v , then $y(v) = x$ for all correct nodes v .
- (3) *Termination.* All correct nodes decide on $y(v)$ and terminate within R rounds.

This modification is straightforward, as the following lemma shows.

LEMMA 30. *Let C be a randomised synchronous consensus routine that terminates in R rounds in expectation and deterministically satisfies agreement and validity conditions. Then there exists a randomised synchronous consensus routine C' that deterministically satisfies validity and terminates within $2R$ rounds, and satisfies agreement with probability at least $1/2$. All other properties, such as message size and resilience, of C and C' are the same.*

PROOF. The modified algorithm operates as follows. We run the original algorithm for (up to) $2R$ rounds. If it terminates at node v , then node v outputs the decision of the algorithm. Otherwise, node v outputs its input value so that $y(v) = x(v)$. This deterministically guarantees validity: if all correct nodes have the same input, the original algorithm can only output that value. Concerning agreement, observe that by Markov's bound, the original algorithm has terminated at all nodes within $2R$ rounds with probability at least $1/2$. Accordingly, agreement holds with probability at least $1/2$. $\square$

We remark that the construction from [24] that generates silent consensus routines out of regular ones also applies to randomised algorithms (as produced by Lemma 30), that is, we can obtain suitable randomised silent consensus routines to be used in our framework.

Our framework makes use of consensus in the construction underlying Theorem 2 only. For stabilisation, we need a constant number of consecutive consensus instances to succeed. Thus, a constant probability of success for each individual consensus instance is sufficient to maintain an expected stabilisation time of $O(R)$ for each individual level in the stabilisation hierarchy. This is summarised in the following variant of Theorem 2.

COROLLARY 9. *Let $f \geq 0$ and $n > 3f$. Suppose, for a network of n nodes, there exist*

- *an f -resilient resynchronisation algorithm B with skew $\rho \in O(d)$ and separation window $\Psi \geq \Psi_0$ for a sufficiently large $\Psi_0 \in O(R)$ and*
- *an f -resilient randomised synchronous consensus algorithm C ,*

where C runs in $R = R(f)$ rounds, lets nodes send at most $M = M(f)$ bits per round and channel, and agreement holds with constant probability. Then there exists a randomised f -resilient pulse synchronisation algorithm A for n nodes with skew $\sigma = 2d$ and accuracy bounds $\Phi^-, \Phi^+ \in \Theta(R)$ that stabilises in expected $T(B) + O(R)$ time and has nodes send $M(B) + O(M)$ bits per time unit and channel.

Note that once one of the underlying pulse synchronisation algorithms used in the construction of Theorem 3 stabilises, the resynchronisation algorithm itself stabilises deterministically, as it does not make use of consensus or randomisation. Applying linearity of expectation and the same recursive pattern as before, Theorem 1 thus generalises as follows:

COROLLARY 10. *Let $\langle C, R, M, N \rangle$ be a family of randomised synchronous consensus routines, where each $C \in C$ satisfies agreement with constant probability. Then, for any $f \geq 0$, $n \geq N(f)$, and $T \geq T_0$*

for some $T_0 \in \Theta(R(f))$, there exists a T -pulser $\mathbf{A}$ with skew $2d$. The number of bits $M(\mathbf{A})$ sent per time unit and channel and the expected stabilisation time $T(\mathbf{A})$ satisfy

$$T(\mathbf{A}) \in O\left(d + \sum_{k=0}^{\lceil \log f \rceil} R(2^k)\right) \quad \text{and} \quad M(\mathbf{A}) \in O\left(1 + \sum_{k=0}^{\lceil \log f \rceil} M(2^k)\right),$$

where the sums are empty when $f = 0$.

However, while randomised algorithms can be more efficient, typically they require additional restrictions on the model, e.g., that the adversary must not be able to predict future random decisions. Naturally, such restrictions then also apply when applying Corollary 9 and, subsequently, Corollary 10. A typical assumption is that communication is via *private channels*. That is, the faulty nodes' behaviour at time t is a function of all communication from correct nodes to faulty nodes during the interval $[0, t]$, the inputs, and the consensus algorithm only.

We can now obtain pulse synchronisation algorithms that are efficient both with respect to stabilisation time and communication. For example, we can make use of the randomised consensus algorithm by King and Saia [21].

THEOREM 6 ([21] AND LEMMA 30). *Suppose communication is via private channels. There is a family of randomised synchronous consensus routines that satisfy the following properties:*

- *the algorithm satisfies agreement with constant probability,*
- *the algorithm satisfies validity,*
- *the algorithm terminates in $R(f) \in \text{polylog } f$ rounds,*
- *the number of bits sent by each node in each round is at most $M(f) \in \text{polylog } f$, and*
- *$n(f) > (3 + \epsilon)f$ for a constant $\epsilon > 0$ that can be freely chosen upfront.*

We point out that the algorithm by King and Saia [21] actually satisfies stronger bounds on the total number of bits sent by each node than what is implied by our statement. As our framework requires nodes to broadcast a constant number of bits per time unit and level of recursion of the construction, we obtain the following corollary.

COROLLARY 2. *Suppose we have private channels. For any $f \geq 0$, constant $\epsilon > 0$, and $n > (3 + \epsilon)f$, there is a randomised f -resilient $(\text{polylog } f)$ -pulser over n nodes that stabilises in $\text{polylog } f$ time w.h.p. and has nodes broadcast $\text{polylog } f$ bits per time unit.*

Note that it is trivial to boost the probability for stabilisation by repetition, as the algorithm must stabilise in $\text{polylog } f$ time regardless of the initial system state. This was exploited in the above corollary. However, in case of a uniform (or slowly growing) running time as function of f , it is useful to apply concentration bounds to show a larger probability of stabilisation. Concretely, the algorithm by Feldman and Micali [18] offers constant expected running time, regardless of f ; this translates to constant probability of success for an $O(1)$ -round algorithm in our setting.

THEOREM 7 ([18] AND LEMMA 30). *Suppose that communication is via private channels. There exists a family of randomised synchronous consensus routines that satisfy the following properties:*

- *the algorithm satisfies agreement with constant probability,*
- *the algorithm satisfies validity,*
- *the algorithm terminates in $R(f) \in O(1)$ rounds,*
- *the total number of bits broadcasted by each node is $\text{poly } f$,*
- *and $n(f) = 3f + 1$.*

Employing this consensus routine, every $O(1)$ time units there is a constant probability that the next level of recursion stabilises. Applying Chernoff's bound over the (at most) $\log f$ recursive levels of stabilisation, this yields stabilisation in $O(\log f)$ time with high probability.

COROLLARY 3. *Suppose we have private channels. For any $f \geq 0$ and $n > 3f$, there is a randomised f -resilient $\Theta(\log f)$ -pulser over n nodes that stabilises in $O(\log f)$ time w.h.p. and has nodes broadcast poly f bits per time unit.*

8 CONCLUSIONS

In this work, we have seen that self-stabilising pulse synchronisation under Byzantine faults can be achieved efficiently in the bounded-delay model with bounded clock drift: the problem reduces to the task of solving (non-stabilising) synchronous binary consensus efficiently. With *deterministic* algorithms, a linear stabilisation time in the number f of faults is possible with nodes broadcasting $O(\log f)$ per time unit. On the other hand, we see that one can obtain sublinear time algorithms by using randomisation at the expense of more bits broadcast per time unit.

We now conclude by highlighting some interesting open problems in the area:

- The construction presented here was based on a reduction *to* consensus. This raises the question whether there is a reduction *from* consensus, that is, is pulse synchronisation at least as hard as consensus? As no reduction in the other direction is known, the true complexity of pulse synchronisation still remains an open question. It may very well be that pulse synchronisation is strictly easier than synchronous consensus.
- The reduction presented in this work is fairly complicated. Are there *simple* and efficient algorithms for achieving pulse synchronisation in a self-stabilising manner?
- Can the techniques used in this work be used to make existing practical non-self-stabilising clock synchronisation algorithms self-stabilising?

ACKNOWLEDGMENTS

We are grateful to Danny Dolev for numerous discussions on the pulse synchronisation problem and detailed comments on early drafts of this article. We also wish to thank Borzoo Bonakdarpour, Janne H. Korhonen, Christian Scheideler, Jukka Suomela, and anonymous reviewers for their helpful comments. Part of this work was done while JR was affiliated with Helsinki Institute for Information Technology HIIT, Department of Computer Science, Aalto University and University of Helsinki.

REFERENCES

- [1] Marcos K. Aguilera and Sam Toueg. 1999. Simple bivalency proof that t -resilient consensus requires $t + 1$ rounds. *Information Processing Letters* 71, 3 (1999), 155–158. DOI: [10.1016/S0020-0190\(99\)00100-3](https://doi.org/10.1016/S0020-0190(99)00100-3)
- [2] Michael Ben-Or. 1983. Another advantage of free choice: Completely asynchronous agreement protocols. In *Proceedings of the 2nd Annual ACM Symposium on Principles of Distributed Computing (PODC'83)*. 27–30. DOI: [10.1145/800221.806707](https://doi.org/10.1145/800221.806707)
- [3] Michael Ben-Or, Danny Dolev, and Ezra N. Hoch. 2008. Fast self-stabilizing Byzantine tolerant digital clock synchronization. In *Proceedings of the 27th ACM Symposium on Principles of Distributed Computing (PODC'08)*. 385–394. DOI: [10.1145/1400751.1400802](https://doi.org/10.1145/1400751.1400802)
- [4] Piotr Berman, Juan A. Garay, and Kenneth J. Perry. 1989. Towards optimal distributed consensus. In *Proceedings of the 30th Annual Symposium on Foundations of Computer Science (FOCS'89)*. 410–415. DOI: [10.1109/SFCS.1989.63511](https://doi.org/10.1109/SFCS.1989.63511)
- [5] Piotr Berman, Juan A. Garay, and Kenneth J. Perry. 1992. Bit optimal distributed consensus. In *Computer Science: Research and Applications*. 313–321. DOI: [10.1007/978-1-4615-3422-8_27](https://doi.org/10.1007/978-1-4615-3422-8_27)
- [6] Ariel Daliot, Danny Dolev, and Hanna Parnas. 2003. Self-stabilizing pulse synchronization inspired by biological pacemaker networks. In *Proceedings of the 6th International Symposium on Self-Stabilizing Systems (SSS'03)*. 32–48. DOI: [10.1007/3-540-45032-7_3](https://doi.org/10.1007/3-540-45032-7_3)

- [7] Edsger W. Dijkstra. 1974. Self-stabilizing systems in spite of distributed control. *Communications of the ACM* 17, 11 (1974), 643–644. DOI: [10.1145/361179.361202](https://doi.org/10.1145/361179.361202)
- [8] Danny Dolev. 1982. The Byzantine generals strike again. *Journal of Algorithms* 3, 1 (1982), 14–30. DOI: [10.1016/0196-6774\(82\)90004-9](https://doi.org/10.1016/0196-6774(82)90004-9)
- [9] Danny Dolev and Ezra N. Hoch. 2007. Byzantine self-stabilizing pulse in a bounded-delay model. In *Proceedings of the 9th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS'07)*. 234–252. DOI: [10.1007/978-3-540-76627-8_19](https://doi.org/10.1007/978-3-540-76627-8_19)
- [10] Danny Dolev and Rüdiger Reischuk. 1985. Bounds on information exchange for Byzantine agreement. *Journal of the ACM* 32, 1 (1985), 191–204. DOI: [10.1145/2455.214112](https://doi.org/10.1145/2455.214112)
- [11] Danny Dolev, Joseph Y. Halpern, and H. Raymond Strong. 1986a. On the possibility and impossibility of achieving clock synchronization. *Journal of Computer and System Sciences* 32, 2 (1986a), 230–250. DOI: [10.1016/0022-0000\(86\)90028-0](https://doi.org/10.1016/0022-0000(86)90028-0)
- [12] Danny Dolev, Nancy A. Lynch, Shlomit S. Pinter, Eugene W. Stark, and William E. Weihl. 1986b. Reaching approximate agreement in the presence of faults. *Journal of the ACM* 33, 3 (1986b), 499–516. DOI: [10.1145/5925.5931](https://doi.org/10.1145/5925.5931)
- [13] Danny Dolev, Matthias Függer, Christoph Lenzen, and Ulrich Schmid. 2014. Fault-tolerant algorithms for tick-generation in asynchronous logic. *Journal of the ACM* 61, 5 (2014), 30:1–30:74. DOI: [10.1145/2560561](https://doi.org/10.1145/2560561)
- [14] Danny Dolev, Keijo Heljanko, Matti Järvisalo, Janne H. Korhonen, Christoph Lenzen, Joel Rybicki, Jukka Suomela, and Siert Wieringa. 2016. Synchronous counting and computational algorithm design. *Journal of Computer and System Sciences* 82, 2 (2016), 310–332. DOI: [10.1016/j.jcss.2015.09.002](https://doi.org/10.1016/j.jcss.2015.09.002)
- [15] Shlomi Dolev. 2000. *Self-Stabilization*. Cambridge, MA.
- [16] Shlomi Dolev and Jennifer L. Welch. 2004. Self-stabilizing clock synchronization in the presence of Byzantine faults. *Journal of the ACM* 51, 5 (2004), 780–799. DOI: [10.1145/1017460.1017463](https://doi.org/10.1145/1017460.1017463)
- [17] Alan D. Fekete. 1990. Asymptotically optimal algorithms for approximate agreement. *Distributed Computing* 4, 1 (1990), 9–29. DOI: [10.1007/BF01783662](https://doi.org/10.1007/BF01783662)
- [18] Peasech Feldman and Silvio Micali. An optimal probabilistic algorithm for synchronous Byzantine agreement. *SIAM Journal on Computing* 26, 4 (n.d.), 873–933. DOI: [10.1137/S0097539790187084](https://doi.org/10.1137/S0097539790187084)
- [19] Michael J. Fischer and Nancy A. Lynch. 1982. A lower bound for the time to assure interactive consistency. *Information Processing Letters* 14, 4 (1982), 183–186. DOI: [10.1016/0020-0190\(82\)90033-3](https://doi.org/10.1016/0020-0190(82)90033-3)
- [20] Pankaj Khanchandani and Christoph Lenzen. 2016. Self-stabilizing Byzantine clock synchronization with optimal precision. In *Proceedings of the 18th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS'16)*. 213–230. DOI: [10.1007/978-3-319-49259-9_18](https://doi.org/10.1007/978-3-319-49259-9_18)
- [21] Valerie King and Jared Saia. 2011. Breaking the $O(n^2)$ bit barrier: Scalable byzantine agreement with an adaptive adversary. *Journal of the ACM* 58, 4 (2011), 1–24. DOI: [10.1145/1989727.1989732](https://doi.org/10.1145/1989727.1989732)
- [22] Leslie Lamport and P. M. Melliar-Smith. 1985. Synchronizing clocks in the presence of faults. *Journal of the ACM* 32, 1 (1985), 52–78. DOI: [10.1145/2455.2457](https://doi.org/10.1145/2455.2457)
- [23] Leslie Lamport, Robert Shostak, and Marshall Pease. 1982. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems* 4, 3 (1982), 382–401. DOI: [10.1145/357172.357176](https://doi.org/10.1145/357172.357176)
- [24] Christoph Lenzen and Joel Rybicki. 2016. Near-optimal self-stabilising counting and firing squads. In *Proceedings of the 18th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS'16)*. 263–280. DOI: [10.1007/978-3-319-49259-9_21](https://doi.org/10.1007/978-3-319-49259-9_21)
- [25] Christoph Lenzen, Matthias Függer, Markus Hofstätter, and Ulrich Schmid. 2013. Efficient construction of global time in SoCs despite arbitrary faults. In *Proceedings of the 16th Euromicro Conference Series on Digital System Design (DSD'13)*. 142–151. DOI: [10.1109/DSD.2013.97](https://doi.org/10.1109/DSD.2013.97)
- [26] Christoph Lenzen, Joel Rybicki, and Jukka Suomela. 2017. Efficient counting with optimal resilience. *SIAM Journal on Computing* 64, 4 (2017), 1473–1500. DOI: [10.1137/16M107877X](https://doi.org/10.1137/16M107877X)
- [27] Jennifer Lundelius and Nancy Lynch. 1984. An upper and lower bound for clock synchronization. *Information and Control* 62, 2–3 (1984), 190–204. DOI: [10.1016/S0019-9958\(84\)80033-9](https://doi.org/10.1016/S0019-9958(84)80033-9)
- [28] Marshall C. Pease, Robert E. Shostak, and Leslie Lamport. 1980. Reaching agreement in the presence of faults. *Journal of the ACM* 27, 2 (1980), 228–234. DOI: [10.1145/322186.322188](https://doi.org/10.1145/322186.322188)
- [29] Michael O. Rabin. 1983. Randomized Byzantine generals. In *Proceedings of the 24th Annual Symposium on Foundations of Computer Science (FOCS'83)*. 403–409. DOI: [10.1109/SFCS.1983.48](https://doi.org/10.1109/SFCS.1983.48)
- [30] Michel Raynal. 2010. *Fault-Tolerant Agreement in Synchronous Message-Passing Systems*. Morgan & Claypool.
- [31] T. K. Srikant and Sam Toueg. 1987. Optimal clock synchronization. *Journal of the ACM* 34, 3 (1987), 626–645. DOI: [10.1145/28869.28876](https://doi.org/10.1145/28869.28876)
- [32] Jennifer L. Welch and Nancy Lynch. 1988. A new fault tolerant algorithm for clock synchronization. *Information and Computation* 77, 1 (1988), 1–36. DOI: [10.1016/0890-5401\(88\)90043-0](https://doi.org/10.1016/0890-5401(88)90043-0)

Received January 2018; revised March 2019; accepted June 2019