Near-Optimal Self-Stabilising Counting and Firing Squads

Christoph Lenzen

Department of Algorithms and Complexity,
Max Planck Institute for Informatics

Joel Rybicki

Helsinki Institute for Information Technology HIIT,
Department of Computer Science, Aalto University

Abstract. Consider a fully-connected synchronous distributed system consisting of n nodes, where up to f nodes may be faulty and every node starts in an arbitrary initial state. In the *synchronous counting* problem, all nodes need to eventually agree on a counter that is increased by one modulo some C in each round. In the *self-stabilising firing squad* problem, the task is to eventually guarantee that all non-faulty nodes have simultaneous responses to external inputs: if a subset of the correct nodes receive an external "go" signal as input, then all correct nodes should agree on a round (in the not-too-distant future) in which to jointly output a "fire" signal. Moreover, no node should generate a "fire" signal without some correct node having previously received a "go" signal as input.

We present a framework reducing both tasks to binary consensus at very small cost: we maintain the resilience of the underlying consensus routine, while the stabilisation time and message size are, up to constant factors, bounded by the sum of the cost of the consensus routine for f faults and recursively applying our scheme to $f' < f/2$ faults. For example, we obtain a deterministic algorithm for self-stabilising Byzantine firing squads with optimal resilience $f < n/3$, asymptotically optimal stabilisation and response time $O(f)$, and message size $O(\log f)$.

As our framework does not restrict the type of consensus routines used, we also obtain efficient randomised solutions, and it is straightforward to adapt our framework to allow for $f < n/2$ omission or $f < n$ crash faults, respectively. Our results resolve various open questions on the two problems, most prominently whether (communication-efficient) self-stabilising Byzantine firing squads or (randomised) sublinear-time solutions for either problem exist.

1 Introduction

The design of distributed systems faces several unique issues related to redundancy and fault-tolerance, timing and synchrony, and the efficient use of communication as a resource [28]. In this work, we give near-optimal solutions to two fundamental distributed synchronisation and coordination tasks: the synchronous counting and the firing squad problems. For both tasks, we devise fast self-stabilising algorithms [16] that are not only communication-efficient, but also tolerate the optimal number of permanently faulty nodes. That is, our algorithms efficiently recover from transient failures that may arbitrarily corrupt the state of the distributed system *and* permanently damage a large number of the nodes.

1.1 Synchronous counting and firing squads

We assume a synchronous message-passing model of distributed computation. The distributed system consists of a fully-connected network of n nodes, where up to f of the nodes may be faulty and the initial state of the system is arbitrary. To model the behaviour of faulty nodes, we consider three typical classes of permanent faults:

- crash (the faulty node stops sending information),
- omission (some or all of the messages sent by the faulty node are lost), and
- Byzantine faults (the faulty node exhibits arbitrary misbehaviour).

Note that even though the communication proceeds in a synchronous fashion, the nodes may have different notions of current time due to the arbitrary initial states. However, many typical distributed protocols assume that the system has either been properly initialised or that the nodes should collectively agree on the rounds in which to perform certain actions. Thus, we are essentially faced with the task of having to agree on a common time in a manner that is both self-stabilising and tolerates permanently faulty behaviour from some of the nodes. To address this issue, we study the synchronous counting and firing squad problems, which are among the most fundamental challenges in fault-tolerant distributed systems.

In the *synchronous counting* problem, all the nodes receive well-separated synchronous clock pulses that designate the start of a new round. The received clock pulses are anonymous, and hence, all correct nodes should eventually stabilise and agree on a round counter that increases consistently by one modulo C. The problem is also known as *digital clock synchronisation*, as all non-faulty nodes essentially have to agree on a shared logical clock. A stabilising execution of such a protocol for $n = 4$, $f = 1$, and $C = 3$ is given below:

Clock										
Node 1	1	1	2	1	0	1	2	0	1	2
Node 2	0	1	2	1	0	1	2	0	1	2
Node 3 (faulty)	*	*	*	*	*	*	*	*	*	*
Node 4	2	0	1	0	0	1	2	0	1	2
			stabilisation				counting			

In the *self-stabilising firing squad problem*, the task is to have all correct nodes eventually stabilise and respond to an external input simultaneously. That is, once stabilised, when a sufficiently large (depending on the type of permanent faults) subset of the correct nodes receive an external "go" signal, then all correct nodes should eventually generate a local "fire" event on the same round. The time taken to react to the "go" signal is called the response time. Note that before stabilisation the nodes may generate spurious firing signals, but after stabilisation no correct node should generate a "fire" event without some correct node having previously

received a "go" signal as input. An execution of such a protocol with $n = 4$, $f = 1$, and response time $R = 5$ is illustrated below:

A firing squad protocol can be used, for example, to agree in a self-stabilising manner on when to initiate a new instance of a non-self-stabilising distributed protocol, as response to internal or external "go" inputs.

1.2 Connections to fault-tolerant consensus

Reaching agreement is perhaps the most intrinsic problem in fault-tolerant distributed computing. It is known that both the synchronous counting [14] and the self-stabilising firing squad problem [13] are closely connected to the well-studied consensus problem [23, 30], where each node is given an input bit and the task is to agree on a common output bit such that if every non-faulty node received the same value as input, then this value must also be the output value. Indeed, the connection is obvious on an intuitive level, as in each task the goal is to agree on a common decision (that is, the output bit, clock value, or whether to generate a firing event).

However, the key difference between the problems lies in self-stabilisation. Typically, the consensus problem is considered in a non-self-stabilising setting with only permanent faults (e.g. $f < n/3$ nodes with arbitrary behaviour), whereas synchronous counting copes with both transient and permanent faults. In fact, it is easy to see that synchronous counting is trivial in a non-self-stabilising setting: if all nodes are initialised with the same clock value, then they can simply locally increment their counters each round without any communication. Furthermore, in a properly initialised system, one can easily reduce the firing squad problem to repeatedly calling a consensus routine [5].

Interestingly, imposing the requirement of self-stabilisation – convergence to correct system behavior from arbitrary initial states – reverses the roles. Solving either the synchronous counting or firing squad problem in a self-stabilising manner also yields a solution to binary consensus, but the converse is not true. In fact, in order to internally or externally trigger a consistent execution of a consensus protocol (or any other non-self-stabilising protocol, for that matter), one first needs a self-stabilising synchronous counting or firing squad algorithm, respectively!

In light of this, the self-stabilising variants of both problems are important generalisations of consensus. While considerable research has been dedicated to both tasks [2, 11, 13–15, 17, 24, 26], our understanding is significantly less developed than for the extensively studied consensus problem. Moreover, it is worth noting that all existing algorithms utilise consensus subroutines [11, 24, 26] or shared coins [2], the latter of which essentially solves consensus as well. Given that both tasks are at least as hard as consensus [14], this seems to be a natural approach. However, it raises the question how much of an overhead must be incurred by such a reduction. In this paper, we subsume and improve upon previous results by providing a generic reduction of synchronous counting and self-stabilising firing squad to binary consensus that incurs very small overheads.

1.3 Contributions

We develop a framework for efficiently transforming *non-self-stabilising* consensus algorithms into *self-stabilising* algorithms for synchronous counting and firing squad problems. In particular, the resulting self-stabilising algorithms have the same resilience as the original consensus algorithms, that is, the resulting algorithms tolerate the same number and type of permanent faults as the original consensus algorithm (e.g. crash, omission, or Byzantine faults).

The construction we give incurs a small overhead compared to time and bit complexity of the consensus routines: the stabilisation time and message size are, up to constant factors, given as the sum of the cost of the consensus routine for f faults and recursively applying our scheme to $f' < f/2$ faults. Finally, our construction can be used in conjunction with both deterministic and randomised consensus algorithms. This allows us to also obtain algorithms for randomised variants of the synchronous counting and firing squad problems.

Our novel framework enables us to address several open problems related to self-stabilising firing squads and synchronous counting. We now give a brief summary of the open problems we solve and the new results obtained using our framework.

Self-stabilising firing squads. In the case of self-stabilising firing squads, Dolev et al. [13] posed the following two open problems:

1. Are there solutions that tolerate either omission or Byzantine (i.e., arbitrary) faults?
2. Are there algorithms using $o(n)$-bit messages only?

We answer both questions in the affirmative by giving algorithms that achieve both properties *simultaneously*. Concretely, our framework implies a deterministic solution for the self-stabilising Byzantine firing squad problem that

- tolerates the optimal number of $f < n/3$ Byzantine faulty nodes,
- uses messages of $O(\log f)$ bits, and
- is guaranteed to stabilise and respond to inputs in linear-in-f communication rounds.

Thus, compared to prior state-of-the-art solutions [13], our algorithm tolerates a much stronger form of faulty behaviour and uses exponentially smaller messages, yet retains asymptotically optimal stabilisation and response time. We also obtain algorithms that tolerate $f < n/2$ omission failures and $f < n$ crash failures while retaining a small message size of $O(\log f)$ bits.

Synchronous counting. We attain novel algorithms for synchronous counting, which is also known as self-stabilising Byzantine fault-tolerant *digital clock synchronisation* [2, 17, 21]. Our new algorithms resolve questions left open by our own prior work [26], namely, whether there exist

1. deterministic linear-time algorithms with optimal resilience and message size $o(\log^2 f)$, or
2. randomised sublinear-time algorithms with small bit complexity.

Again, we answer both questions positively using our framework developed in this paper. For the first question, we give linear-time deterministic algorithms that have message size $O(\log f)$ bits. For the second question, we show that our framework can utilise efficient randomised consensus algorithms to obtain probabilistic variants of the synchronous counting and firing squad problems. For example, the result of King and Saia [22] implies algorithms that stabilise with high probability in polylog n rounds and use message size polylog n, assuming private communication links and an adaptive Byzantine adversary corrupting $f < n/(3 + \varepsilon)$ nodes for an arbitrarily small constant $\varepsilon > 0$.

1.4   Related work

In this section, we overview prior work on the synchronous counting and firing squad problems. By now it has been established that both problems [13, 14] are closely connected to the well-studied (non-self-stabilising) consensus [23, 30]. As there exists a vast body of literature on synchronous consensus, we refer the interested reader to e.g. the survey by Raynal [31]. We note that self-stabilising variants of consensus have been studied [1, 8, 9, 18] but in different models of computation and/or failures than what we consider in this work.

Synchronous counting and digital clock synchronisation.   In the past two decades, there has been increased interest in combining self-stabilisation with Byzantine fault-tolerance. One reason is that algorithms in this fault model are very attractive in terms of designing highly-resilient hardware [14]. A large body of work on synchronous counting (also known as digital clock synchronisation) exists [2, 11, 15, 17, 21, 27], comprising both positive and negative results.

In terms of lower bounds, many impossibility results for consensus [10, 12, 20, 30] also directly apply to synchronous counting, as synchronous counting solves binary consensus [14, 15]. In particular, no algorithm can tolerate more than $f < n/3$ Byzantine faulty nodes [30] (unless cryptographic assumptions are made) and any deterministic algorithm needs at least $f + 1$ rounds to stabilise [20].

In a seminal work, Dolev and Welch [17] showed that the task can be solved in a self-stabilising manner in the presence of (the optimal number of) $f < n/3$ Byzantine faults using randomisation; see also [16, Ch. 6]. While this algorithm can be implemented using only constant-size messages, the expected stabilisation time is exponential. Later, Ben-Or et al. [2] showed that it is possible to obtain optimally-resilient solutions that stabilise in expected constant time. However, their algorithm relies on shared coins, which are costly to implement and assume private communication channels.

In addition to the lower bound results, there also exist deterministic algorithms for the synchronous counting problem [11, 15, 21, 27]. Many of these algorithms utilise consensus routines [11, 21, 27], but obtaining fast and communication-efficient solutions with optimal resilience has been a challenge. For example, Dolev and Hoch [11] apply a pipelining technique, where $\Omega(n)$ consensus instances are run in parallel. While this approach attains optimal resilience and linear stabilisation time in f, the large number of parallel consensus instances necessitates large messages.

In order to achieve better communication and state complexity, the use of computational algorithm design and synthesis techniques have also been investigated [4, 15]. While this line of research has produced novel optimal and computer-verified algorithms, so far the techniques have not scaled beyond $f = 1$ faulty node due to the inherent combinatorial explosion in the search space of potential algorithms.

Recently, we gave recursive constructions that achieve linear stabilisation time using only polylogarithmic message size and state bits per node [24, 26]; see also the extended and revised version [27]. However, our previous constructions relied on specific (deterministic) consensus routines and their properties in a relatively ad hoc manner. In contrast, our new framework presented here lends itself to any (possibly randomised) synchronous consensus routine and improves the best known upper bound on the message size to $O(\log f)$ bits. Currently, it is unknown whether it is possible to deterministically achieve message size $o(\log f)$.

Firing squads.   In the original formulation of the firing squad synchronisation problem, the system consists of an n-length path consisting of finite state machines (whose number of states is independent of n) and the goal is to have all machines switch to the same "fire" state simultaneously after one node receives a "go" signal. This formulation of the problem has been attributed to John Myhill and Edward Moore and has subsequently been studied in various settings; see e.g. [29] for survey of early work related to the problem.

In the distributed computing community, the firing squad problem has been studied in fully-conneted networks in the presence of faulty nodes. Similarly to synchronous counting, the firing squad problem is closely connected to Byzantine agreement and simultaneous consensus [5–7, 13, 19]. Both Burns and Lynch [5] and Coan et al. [6] studied the firing squad problem in the context of Byzantine failures. Burns and Lynch [5] considered both permissive and strict variants of the problem (i.e., whether faulty nodes can trigger a firing event or not) and showed that both can be solved using Byzantine consensus algorithms with only a relatively small additional overhead in the number of communication rounds and total number of bits communicated. On the other hand, Coan et al. [6] gave *authenticated* firing squad algorithms for various Byzantine fault models. Coan and Dwork [7] gave time lower bounds of $f + 1$ rounds for deterministic and randomised algorithms solving the firing squad problem in the crash fault model.

However, neither the solutions of Burns and Lynch [5] or Coan et al. [6] are self-stabilising or use small messages. Almost two decades later, Dolev et al. [13] gave the first self-stabilising algorithm for the firing squad problem. In particular, their solution has optimal stabilisation time and response time depending on the fault pattern. However, their algorithm tolerates only crash faults and uses messages of size $\Theta(n \log n)$ bits. In this work, we improve on this result by achieving Byzantine fault-tolerance using messages of $O(\log n)$ bits.

1.5 Outline of the paper

The paper is structured as follows. For the sake of brevity, in the main body of the paper we confine the presentation to Byzantine faults; crash and omission faults are discussed in Appendix A. In Section 2, we give formal definitions related to the model of computation, synchronous counting, and firing squads. In the sections following this, we show our main result in a top-down fashion. We introduce a series of new problems and give reductions between them:

- Section 3 shows how to obtain synchronous counting and firing squad algorithms that rely on binary consensus routines and *strong pulsers*,
- Section 4 devises strong pulsers with the help of *weak pulsers* and multivalued consensus,
- Section 5 constructs weak pulsers using *silent consensus* and less resilient strong pulsers.

Finally, Section 6 combines the results of Section 4 and Section 5 to obtain a recursive construction for strong pulsers used by the algorithms given in Section 3. We conclude the paper with a brief discussion of additional results implied by our framework.

2 Preliminaries

In this section, we first fix some basic notation, then describe the model of computation and give formal definitions of the synchronous counting, self-stabilising firing squad, and consensus problems.

2.1 Notation

We use $\mathbb{N} = \{1, 2, \ldots\}$ to denote the set of positive integers and $\mathbb{N}_0 = \{0\} \cup \mathbb{N}$ to denote the set of all non-negative integers. For any $k \in \mathbb{N}$, we write $[k] = \{0, \ldots, k-1\}$ to be the set of first k non-negative integers.

2.2 Model of computation

We consider a fully-connected synchronous network on node set V consisting of $n = |V|$ processors. We assume there exists a subset of $F \subseteq V$ faulty nodes that is (at least initially) unknown to all nodes, where the upper bound f on the size $|F| \leq f$ is known by the nodes. We say that nodes in $V \setminus F$ are correct and nodes in F are faulty.

All correct nodes in the system will follow a given algorithm **A** that is the same for all the nodes in the system. The execution proceeds in synchronous rounds, where in each round $t \in \mathbb{N}$ the nodes take the following actions in lock-step:

1. perform local computations,
2. send messages to other nodes, and
3. receive messages from other nodes.

We assume nodes can distinguish the sender of incoming messages, e.g., by having a port numbering assigning distinct labels to the incoming links from each other node. In particular, our construction does not require any port numbering of the outgoing links or unique identifiers, unless the underlying consensus routine does. We say that an algorithm **A** has message size $M(\mathbf{A})$ if no node sends more than $M(\mathbf{A})$ bits to any other node during a single round.

The local computations of a node v determine the decision which messages to send to other nodes and what is the new state of the node v. As we are interested in self-stabilising algorithms, the initial state of a node is arbitrary; this is equivalent to assuming that transient faults have arbitrarily corrupted the state of each node, but the transient faults have ceased by the beginning of the first round.

As mentioned above, we allow for additional (possibly permanent) *Byzantine* faults. A Byzantine faulty node $v \in F$ may deviate from the algorithm arbitrarily, i.e., send arbitrary messages in each round. In particular, a Byzantine faulty node can send different messages to each correct node in the system, even if the algorithm specifies otherwise. Since we consider deterministic algorithms, the meaning of "arbitrary" in this context is that the algorithm must succeed for any possible choice of behavior of the faulty nodes. We require that $f = |F| < n/3$, as otherwise none of the problems we consider can be solved due to the impossibility of consensus under $f \geq n/3$ Byzantine faults [30].

2.3 Synchronous counting

In the *synchronous C-counting* problem, the task is to have each node $v \in V$ output a counter value $c(v, t) \in [C]$ on each round $t \in \mathbb{N}$ in a consistent manner. We say that an execution of an algorithm *stabilises in round t* if and only if all $t \leq t' \in \mathbb{N}$ and $v, w \in V \setminus F$ satisfy

- **Agreement:** $c(v, t') = c(w, t')$ and
- **Consistency:** $c(v, t' + 1) = c(v, t') + 1 \bmod C$.

We say that **A** is an f-resilient C-counting algorithm that stabilises in time t if all executions with at most f faulty nodes stabilise by round t. The stabilisation time $T(\mathbf{A})$ of **A** is the maximum such t over all executions.

2.4 Self-stabilising firing squad

In the self-stabilising Byzantine firing squad problem, in each round $t \in \mathbb{N}$, each node $v \in V$ receives an external input $\mathrm{GO}(v, t) \in \{0, 1\}$. Moreover, the algorithm determines an output $\mathrm{FIRE}(v, t) \in \{0, 1\}$ at each node $v \in V$ in each round $t \in \mathbb{N}$. We say that an execution of an algorithm *stabilises in round* $t \in \mathbb{N}$ if the following three properties hold:

- **Agreement:** $\mathrm{FIRE}(v, t') = \mathrm{FIRE}(w, t')$ for all $v, w \in V \setminus F$ and $t \leq t' \in \mathbb{N}$.
- **Safety:** If $\mathrm{FIRE}(v, t_F) = 1$ for $v \in V \setminus F$ and $t \leq t_F \in \mathbb{N}$, then there is $t_F \geq t_G \in \mathbb{N}$ such that

  (i) $\mathrm{GO}(w, t_G) = 1$ for some $w \in V \setminus F$ and
  (ii) $\mathrm{FIRE}(v, t') = 0$ for all $t' \in \{t_G + 1, \ldots, t_F - 1\}$.

- **Liveness:** If $GO(v, t_G) = 1$ for at least $f + 1$ nodes $v \in V \setminus F$ and $t \le t_G \in \mathbb{N}$, then $FIRE(v, t_F) = 1$ for all nodes $v \in V \setminus F$ and some $t_G < t_F \in \mathbb{N}$.

Note that the liveness condition requires $f + 1$ correct nodes to observe a GO input, as otherwise it would be impossible to guarantee that a correct node observed a GO input when firing; this corresponds to the definition of a strict Byzantine firing squad [5]. We say that an execution stabilised by round t has *response time R from round t* if

(i) when firing is required in response to (sufficiently many) GO inputs of 1 in round $t_G \ge t$, this happens no later than round $t_G + R$, and
(ii) when the squad fires in round $t_F \ge t$, there was sufficient support (in terms of GO inputs of 1) justifying this in a round t_G with $t_F > t_G \ge t_F - R$.

Finally, we say that an algorithm $\mathbf{F}$ is an f-resilient firing squad algorithm with stabilisation time $T(\mathbf{F})$ and response time $R(\mathbf{F})$ if in any execution of the system with at most f faulty nodes there is a round $t \le T(\mathbf{F})$ such that the algorithm stabilised and has response time at most $R(\mathbf{F})$ from round t.

2.5 Consensus

Let us conclude this section by definining the *multivalued consensus problem*. Unlike the synchronous counting and self-stabilising firing squad problems, the standard definition of consensus does *not* require self-stabilisation: we assume that all nodes start from a fixed starting state and the algorithm terminates in finitely many communication rounds.

In the multivalued consensus problem for $L > 1$ values, each node $v \in V$ receives an input value $x(v) \in [L]$ and the task is to have all correct nodes output the same value $y \in [L]$. We say that an algorithm $\mathbf{C}$ is an f-resilient $T(\mathbf{C})$-round consensus algorithm if the following conditions hold when there are at most f faulty nodes:

- **Termination:** Each $v \in V \setminus F$ decides on an output $y(v) \in [L]$ by the end of round $T(\mathbf{C})$.
- **Agreement:** For all $v, w \in V \setminus F$, it holds that $y(v) = y(w)$.
- **Validity:** If there exists $x \in [L]$ such that for all $v \in V \setminus F$ it holds that $x(v) = x$, then each $v \in V \setminus F$ outputs the value $y(v) = x$.

We remark that one may ask for stronger validity conditions, but for our purposes this condition is sufficient. The *binary consensus* problem is the special case of $L = 2$ of the above multivalued consensus problem. In the case of binary consensus, the stated validity condition is equivalent to requiring that if $v \in V \setminus F$ outputs $y(v) = x \in \{0, 1\}$, then some $w \in V \setminus F$ has input value $x(w) = x$.

Later, we utilise the fact that multivalued consensus can be reduced to binary consensus with only a small overhead in time. In [25], it is shown how to do this with 1-bit messages and an additive overhead of $O(\log L)$ rounds, preserving resilience.

Theorem 1 ([25]). *Let $L > 1$. Given an f-resilient binary consensus algorithm $\mathbf{C}$, we can solve L-value consensus in $O(\log L) + T(\mathbf{C})$ rounds using $M(\mathbf{C})$-bit messages while tolerating f faults.*

3 Synchronous counting and firing squads

In this section, we give a firing squad algorithm with asymptotically optimal stabilisation and response times. The algorithm relies on two auxiliary routines: a so-called strong pulser and a consensus algorithm. We start with a discussion on strong pulsers.
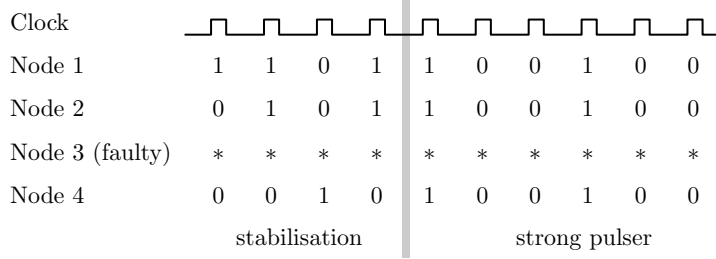
Clock										
Node 1	1	1	0	1	1	0	0	1	0	0
Node 2	0	1	0	1	1	0	0	1	0	0
Node 3 (faulty)	*	*	*	*	*	*	*	*	*	*
Node 4	0	0	1	0	1	0	0	1	0	0
		stabilisation					strong pulser			

Figure 1: An example execution of a strong 3-pulser on $n = 4$ nodes with $f = 1$ faulty node.

3.1 Strong pulsers and counting

Our approach to the firing squad problem is to solve it by repeated consensus, which in turn is controlled by a joint counter. To minimise message size, however, we will not communicate counter values directly. Instead we make use of what we call a *strong pulser*.

Definition 1 (Strong pulser). *An algorithm* **P** *is an f-resilient strong Ψ-pulser that stabilises in $T(\mathbf{P})$ rounds if it satisfies the following conditions in the presence of at most f faulty nodes. Each node $v \in V$ produces an output bit $p(v,t) \in \{0,1\}$ on each round $t \in \mathbb{N}$. We say that v generates a pulse in round t if $p(v,t) = 1$ holds. We require that there is a round $t_0 \leq T(\mathbf{P})$ such that:*

S1. For any $v \in V \setminus F$ and round $t = t_0 + k\Psi$, where $k \in \mathbb{N}_0$, it holds that $p(v,t) = 1$.
S2. For any $v \in V \setminus F$ and round $t \geq t_0$ satisfying $t \neq t_0 + k\Psi$ for $k \in \mathbb{N}_0$, we have $p(v,t) = 0$.

Put otherwise, a strong Ψ-pulser consistently generates pulses at all non-faulty nodes exactly every Ψ rounds. Figure 1 illustrates an execution of a strong pulser with $\Psi = 3$. It is straightforward to see that strong pulsers and synchronous counting are almost equivalent.

Lemma 1. *Let $C \in \mathbb{N}$ and $\Psi \in \mathbb{N}$. If C divides Ψ, then a strong Ψ-pulser that stabilises in T rounds implies a synchronous C-counter that stabilises in at most T rounds. If Ψ divides C, then a synchronous C-counter that stabilises in T rounds implies a strong Ψ-pulser that stabilises in at most $T + \Psi - 1$ rounds.*

Proof. For the first claim, set $c(v,t) = 0$ in any round t for which $p(v,t) = 1$ and $c(v,t) = c(v,t-1) + 1 \bmod C$ in all other rounds. For the second claim, set $p(v,t) = 1$ in all rounds t in which $c(v,t) \bmod \Psi = 0$ and $p(v,t) = 0$ in all other rounds. $\square$

Another way of interpreting this relation is to view a strong Ψ-pulser as a different encoding of the output of a Ψ-counter: since the system is synchronous, it suffices to communicate when the counter overflows to value 0 and otherwise count locally. This saves bandwidth when communicating the state of the counter.

3.2 Firing squads via pulsers and consensus

We now show how an f-resilient strong pulser and a f-resilient binary consensus algorithm can be used to devise an f-resilient firing squad algorithm. As a strong pulser can be used to control repeated execution of a non-self-stabilising algorithm, it enables us to run consensus on whether a firing event should be triggered or not repeatedly. As the firing squad problem is at least as hard as consensus, this maintains asymptotically optimal round complexity.

Recall that for the *Byzantine* firing squad problem, we are interested in a liveness condition in which a firing event needs to be generated if at least $f + 1$ non-faulty nodes $v \in V \setminus F$ recently saw $\mathrm{GO}(v,t) = 1$ on some round t. To this end, we have each node continuously inform all other

nodes about its GO values (i.e. their received input signals). Whenever node $v \in V$ sees $f + 1$ nodes $w \in V$ claim $\text{GO}(w, t) = 1$, it will memorise this and use input $x(v) = 1$ for the next consensus instance. Otherwise, it will use the input value $x(v) = 0$; this ensures that at least one non-faulty node w had $\text{GO}(w, t) = 1$ recently in case v uses input $x(v) = 1$. The validity condition of the (arbitrary) $T(\mathbf{C})$-round consensus routine $\mathbf{C}$ thus ensures both liveness and safety for the resulting firing squad algorithm. Apart from $\mathbf{C}$, the algorithm concurrently runs a strong Ψ-pulser $\mathbf{P}$ for some $\Psi > T(\mathbf{C})$.

The firing squad algorithm. Given a strong Ψ-pulser algorithm $\mathbf{P}$ and a binary consensus algorithm $\mathbf{C}$, each node v stores the following variables on every round t:

- $p(v, t) \in \{0, 1\}$, the output variable of $\mathbf{P}$,
- $x(v, t) \in \{0, 1\}$ and $y(v, t) \in \{0, 1\}$, the input and output variables of $\mathbf{C}$, and
- $m(v, t) \in \{0, 1\}$, an auxiliary variable used to memorise whether sufficiently many GO signals were received to warrant a firing event.

The algorithm consists of each node v executing the following operations[1] in each round $t \in \mathbb{N}$:

1. Broadcast $\text{GO}(v, t)$.
2. If received at least $f + 1$ times $\text{GO}(w, t - 1) = 1$, then set $x(v, t) = 1$ and $m(v, t) = 1$. Otherwise, set $x(v, t) = x(v, t - 1)$ and $m(v, t) = m(v, t - 1)$.
3. If $p(v, t) = 1$, start executing a new instance of $\mathbf{C}$ using the value $x(v, t - 1)$ as input and set $m(v, t) = 0$ while aborting any previously running instance. More specifically, this entails the following:

   - Maintain a local round counter r, which is initialised to 1 on round t and increased by 1 after each round.
   - Maintain the local state variables related to the consensus routine $\mathbf{C}$.
   - On each round, execute round r of algorithm $\mathbf{C}$; if the state variables indicate that $\mathbf{C}$ terminated at v, then do nothing.
   - At the end of the round when r becomes $T(\mathbf{C})$, stop the simulation (indicating this, e.g., by setting $r(v) = \bot$) and locally output the value of $y(v)$ computed by the simulation of $\mathbf{C}$.

4. If $\mathbf{C}$ outputs $y(v, t) = 1$ on round t, then output $\text{FIRE}(v, t) = 1$ and set $x(v, t) = 0$. Otherwise, set $\text{FIRE}(v, t) = 0$.
5. If $\mathbf{C}$ outputs $y(v, t) = 0$ on round t and $m(v, t) = 0$, then set $x(v, t) = 0$.

We now show that the above algorithm satisfies the properties required from a self-stabilising firing squad.

Theorem 2. *Suppose there exists an f-resilient strong Ψ-pulser $\mathbf{P}$ and a consensus algorithm $\mathbf{C}$, where $\Psi > T(\mathbf{C})$. Then there exists an f-resilient firing squad algorithm $\mathbf{F}$ that*

- *stabilises in time $T(\mathbf{F}) \le T(\mathbf{P}) + T(\mathbf{C}) + 1$,*
- *has response time $R(\mathbf{F}) \le \Psi + T(\mathbf{C})$, and*
- *uses message of size $M(\mathbf{F}) \le M(\mathbf{P}) + M(\mathbf{C}) + 1$ bits.*

Proof. Let $\mathbf{F}$ be the algorithm described above. We now argue that the algorithm satisfies the three properties given in Section 2.4: (i) agreement, (ii) safety, and (iii) liveness.

(i) Denote by $t_0 \le T(\mathbf{P})$ the round in which the execution of the strong Ψ-pulser $\mathbf{P}$ has stabilised and generated a pulse. That is, for rounds $t \ge t_0$ we have that $p(v, t) = 1$ is equivalent

---

[1]For better readability, we allow for statements about what a node communicates to appear anywhere in the description. Note, however, that sending operations happen after local computation, i.e., only information sent in the previous rounds is available for computations.

to $t = t_0 + k\Psi$ for some $k \in \mathbb{N}_0$. This implies that the algorithm will correctly simulate instances of the consensus routine $\mathbf{C}$ and locally output its decision on rounds $t_k = t_0 + T(\mathbf{C}) + k\Psi$ for $k \in \mathbb{N}_0$. The agreement property of the firing squad thus follows from the agreement property of consensus for all rounds $t \geq t_0$, as $\mathrm{FIRE}(v,t) = 1$ if and only if $t = t_k$ and the simulation of $\mathbf{C}$ output the value $y(v,t) = 1$ in Step 4.

(ii) Concerning safety, suppose $v \in V \setminus F$ outputs $\mathrm{FIRE}(v,t_F) = 1$ in round $t_F > t_0 + T(\mathbf{C})$. By the above discussion and the validity property of consensus, this implies that there was some node $w \in V \setminus F$ that started a (successfully and completely simulated) instance of $\mathbf{C}$ with input $x(w,t_k) = 1$ in round $t_k = t_F - T(\mathbf{C}) = t_0 + k\Psi$ for some $k \in \mathbb{N}_0$. Assume for contradiction that there are no $u \in V \setminus F$ and $t_k - \Psi \leq t_G < t_k$ satisfying $\mathrm{GO}(u,t_G) = 1$. Then, w does not set $x(w,t')$ or $m(w,t')$ to 1 in rounds $t' \in \{t_k - \Psi + 1, \ldots, t_k\}$ in Step 2. However, in round $t_F - T(\mathbf{C}) - \Psi = t_0 + (k-1)\Psi = t_{k-1}$ node w set $m(w,t_{k-1}) = 0$ and thus w sets $x(w,t_F - \Psi) = 0$ later in round $t_F - \Psi$, the round in which the previous instance of $\mathbf{C}$ locally output some value (Steps 4 and 5). This contradicts the fact that $x(w,t_k) = 1$ in round $t_k = t_F - T(\mathbf{C}) > t_F - \Psi$. Hence, there are $u \in V \setminus F$ and $t_F - T(\mathbf{C}) - \Psi \leq t_G < t_k$ such that $\mathrm{GO}(u,t_G) = 1$.

The safety property requires also that $\mathrm{FIRE}(v,t') = 0$ for all $t' \in \{t_G + 1, \ldots, t_F - 1\}$. The time t_G shown to exist by the above reasoning does not satisfy this if and only if some instance of $\mathbf{C}$ locally outputs $y(v,t') = 1$ at node v in such a round t'. The only possible such round t' is $t_F - \Psi$, as $t_F - 2\Psi < t_F - T(\mathbf{C}) - \Psi$. However, in this case, each $w \in V \setminus F$ sets $x(w,t_F - \Psi) = 0$ in round $t_F - \Psi$ regardless of $m(w,t_F - \Psi)$ in Step 4, and we can conclude that some $w \in V \setminus F$ must set $x(w,t'') = 1$ in some round $t'' \in \{t_F - \Psi + 1, \ldots, t_F - T(\mathbf{C})\}$. As above, it follows that there is a round $t_G \in \{t_F - \Psi, \ldots, t_F - T(\mathbf{C})\}$ and a node $u \in V \setminus F$ such that $\mathrm{GO}(u,t_G) = 1$. Overall, we see that the safety condition for a firing squad algorithm with response time $\Psi + T(\mathbf{C})$ is met in rounds $t_F \geq t_0 + T(\mathbf{C}) + 1$.

(iii) It remains to address liveness. Suppose at least $f + 1$ nodes $v \in V \setminus F$ satisfy $\mathrm{GO}(v,t_G) = 1$ in some round $t_G \geq t_0 - 1$. Then, in round $t_G + 1 \geq t_0$ all nodes $v \in V \setminus F$ set $x(v,t_G + 1) = 1$ and $m(v,t_G + 1) = 1$ according to Step 2. Assume for contradiction that $\mathrm{FIRE}(v,t) = 0$ for all $t \in \{t_G + 1, \ldots, t_G + T(\mathbf{C}) + \Psi\}$. Denote by $t_G + 1 \leq t_k \leq t_G + \Psi$ the unique round such that $t_k = t_0 + k\Psi$ for some $k \in \mathbb{N}_0$. The instance of $\mathbf{C}$ started in this round will satisfy that all correct nodes $v \in V \setminus F$ have input $x(v,t_k) = 1$: by our assumption for contradiction, no node can locally output $y(v,t') = 1$ during rounds $t' \in \{t_G + 1, \ldots, t_G + T(\mathbf{C}) + \Psi\}$; thus, no node can set $x(v,\cdot)$ to 0 without setting $m(v,\cdot)$ to 0 first, which in turn entails to start an instance of $\mathbf{C}$ with the current value of $x(v,\cdot) = 1$ first. By the properties of $\mathbf{C}$, it follows that each $v \in V \setminus F$ locally outputs 1 in round $t_G + 1 < t_k + T(\mathbf{C}) \leq t_G + \Psi + T(\mathbf{C})$, contradicting our previous assumption.

We conclude that our algorithm satisfies the liveness property with response time $\Psi + T(\mathbf{C})$ for rounds $t_G \geq t_0 - 1$. Recalling that $t_0 \leq T(\mathbf{P})$, we see that the algorithm is a firing squad with response time at most $\Psi + T(\mathbf{C})$ that stabilises in $\max\{T(\mathbf{P}), T(\mathbf{P}) + T(\mathbf{C}) + 1, T(\mathbf{P}) - 1\} = T(\mathbf{P}) + T(\mathbf{C}) + 1$ rounds. The bound on the message size follows from the fact that the algorithm $\mathbf{F}$ only broadcasts 1 bit in Step 1 in addition to the messages related to $\mathbf{P}$ and $\mathbf{C}$. $\qquad\square$

4   From weak pulsers to strong pulsers

In the previous section, we established that it suffices to construct suitable strong pulsers to solve the synchronous counting and firing squad problems. We will now reduce the construction of strong pulsers to constructing *weak pulsers.*

4.1   Weak pulsers

A weak Φ-pulser is similar to a strong pulser, but does not guarantee a fixed frequency of pulses. However, it guarantees to *eventually* generate a pulse followed by $\Phi - 1$ rounds of silence.
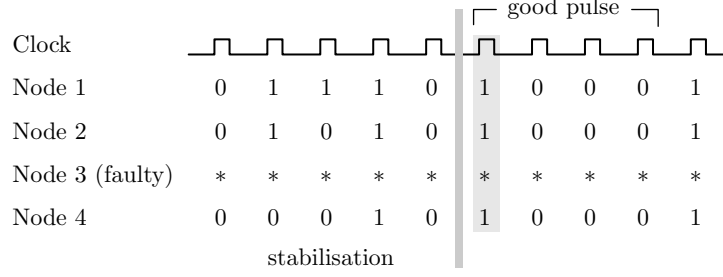
good pulse

Clock										
Node 1	0	1	1	1	0	1	0	0	0	1
Node 2	0	1	0	1	0	1	0	0	0	1
Node 3 (faulty)	*	*	*	*	*	*	*	*	*	*
Node 4	0	0	0	1	0	1	0	0	0	1

stabilisation

Figure 2: An example execution of a weak 4-pulser on $n = 4$ nodes with $f = 1$ faulty node. Eventually, a good pulse is generated, which is highlighted. A good pulse is followed by three rounds in which no correct node generates pulse. In contrast, the pulse two rounds earlier is not good, as it is followed by only one round of silence.

Formally, we define weak pulsers as follows.

Definition 2 (Weak pulsers). *An algorithm* $\mathbf{W}$ *is an f-resilient weak Φ-pulser that stabilises in* $T(\mathbf{W})$ *rounds if the following holds. In each round $t \in \mathbb{N}$, each node $v \in V$ produces an output* $a(v, t)$. *Moreover, there exists a round $t_0 \leq T(\mathbf{W})$ such that*

W1. for all $v, w \in V \setminus F$ and all rounds $t \geq t_0$, $a(v, t) = a(w, t)$,
W2. $a(v, t_0) = 1$ for all $v \in V \setminus F$, and
W3. $a(v, t) = 0$ for all $v \in V \setminus F$ and $t \in \{t_0 + 1, \ldots, t_0 + \Phi - 1\}$.

We say that on round t_0 a good *pulse is generated by* $\mathbf{W}$.

Figure 2 illustrates a weak 4-pulser. Note that while the definition formally only asks for one good pulse, the fact that the algorithm guarantees this property for any starting state implies that there is a good pulse at least every $T(\mathbf{W})$ rounds.

4.2 Constructing strong pulsers from weak pulsers

Recall that a strong pulser can be obtained by having nodes locally count down the rounds until the next pulse, provided we have a way of ensuring that the local counters eventually agree. This can be achieved by using a weak pulser to control a suitable consensus routine, where again we always have only a single instance running at any time. While some instances will be aborted before they can complete, this will not affect the counters, as we only adjust them when the consensus routine completes. On the other hand, the weak pulser guarantees that within $T(\mathbf{W})$ rounds, there will be a pulse followed by $\Phi - 1$ rounds of silence, enabling to complete a run of any consensus routine $\mathbf{C}$ satisfying $T(\mathbf{C}) \leq \Phi$. Thus, for constructing a strong Ψ-pulser, we assume that we have the following f-resilient algorithms available:

- a $T(\mathbf{C})$-round Ψ-value consensus algorithm $\mathbf{C}$ and
- a weak Φ-pulser $\mathbf{W}$ for $\Phi \geq T(\mathbf{C})$.

Given the above two algorithms, we show how to construct an f-resilient strong Ψ-pulser for any $\Psi > 1$. The pulser will stabilise in time $T(\mathbf{W}) + T(\mathbf{C}) + \Psi$ and the message size of the strong pulser will be bounded by $M(\mathbf{W}) + M(\mathbf{C})$.

As mentioned earlier, the idea is have nodes simply count locally between pulses and use the weak pulser to execute a single instance of the consensus algorithm $\mathbf{C}$. Eventually, a good pulse will run an instance consistently and establish agreement among the local counters. Leveraging validity, we can ensure that the counters will never be affected by the consensus routine running in the background again.

11

Variables. Beside the variables of the weak pulser $\mathbf{W}$ and (a single copy of) $\mathbf{C}$, our construction of a strong Ψ-pulser uses the following local variables:

- $a(v,t) \in \{0,1\}$ is the output variable of the weak Φ-pulser $\mathbf{W}$,
- $b(v,t) \in \{0,1\}$ is the output variable of the strong Ψ-pulser we are constructing,
- $c(v,t) \in [\Psi]$ is the local counter keeping track on when the next pulse occurs, and
- $d(v,t) \in \{1,\ldots,T(\mathbf{C})\} \cup \{\bot\}$ keeps track of how many rounds an instance of $\mathbf{C}$ has been executed since the last pulse from the weak pulser $\mathbf{W}$. The value $\bot$ denotes that the consensus routine has stopped.

Strong pulser algorithm. The algorithm is as follows. Each node v executes the weak Φ-pulser algorithm $\mathbf{W}$ in addition to the following instructions on each round t:

1. If $c(v,t) = 0$, then set $b(v,t+1) = 1$ and otherwise $b(v,t+1) = 0$.
2. Set $c'(v,t) = c(v,t)$.
3. If $d(v,t) \neq \bot$, then
   - (a) Execute the instructions of $\mathbf{C}$ for round $d(v,t)$.
   - (b) If $d(v,t) \neq T(\mathbf{C})$, set $d(v,t+1) = d(v,t)+1$.
   - (c) If $d(v,t) = T(\mathbf{C})$, then
     - i. Set $c'(v,t) = y(v,t) + T(\mathbf{C}) \bmod \Psi$, where $y(v,t)$ is the output value of $\mathbf{C}$.
     - ii. Set $d(v,t+1) = \bot$.
4. Update $c(v,t+1) = c'(v,t) + 1 \bmod \Psi$.
5. If $a(v,t) = 1$, then
   - (a) Start a new instance of $\mathbf{C}$ using $c'(v,t)$ as input (resetting all state variables of $\mathbf{C}$).
   - (b) Set $d(v,t+1) = 1$.

In the above algorithm, the first step simply translates the counter value to the output of the strong pulser. We then use a temporary variable $c'(v,t)$ to hold the counter value, which is overwritten by the output of $\mathbf{C}$ (increased by $T(\mathbf{C}) \bmod \Psi$) if it completes a run in this round. In either case, the counter value needs to be increased by $1 \bmod \Psi$ for the next round. The remaining code does the bookkeeping for an ongoing run of $\mathbf{C}$ and starting a new run if the weak pulser generates a pulse.

Observe that in the above algorithm, each node only sends messages related to the weak pulser $\mathbf{W}$ and the consensus algorithm $\mathbf{C}$. Thus, there is no additional overhead in communication and the message size is bounded by $M(\mathbf{W}) + M(\mathbf{C})$. Hence, it remains to show that the local counters $c(v,t)$ implement a strong Ψ-counter.

Theorem 3. *The variables $c(v,t)$ in the above algorithm implement a synchronous Ψ-counter that stabilises in $T(\mathbf{W}) + T(\mathbf{C}) + 1$ rounds and uses messages of at most $M(\mathbf{W}) + M(\mathbf{C})$ bits.*

Proof. Suppose round $t_0 \leq T(\mathbf{W})$ is as in Definition 2, that is, $a(v,t) = a(w,t)$ for all $t \geq t_0$, and a good pulse is generated in round t_0. Thus, all correct nodes participate in simulating an instance of $\mathbf{C}$ during rounds $t_0 + 1, \ldots, t_0 + T(\mathbf{C})$, since no pulse is generated during rounds $t_0 + 1, \ldots, t_0 + T(\mathbf{C}) - 1$, and thus, also no new instance is started in the last step of the code during these rounds.

By the agreement property of the consensus routine, it follows that $c'(v, t_0 + T(\mathbf{C})) = c'(w, t_0 + T(\mathbf{C}))$ for all $v,w \in V \setminus F$ after Step 3ci. By Steps 2 and 4, the same will hold for both $c(\cdot, t')$ and $c'(\cdot, t')$, $t' > t_0 + T(\mathbf{C})$, provided that we can show that in rounds $t' > t$, Step 3ci never sets $c'(v,t)$ to a value different than $c(v,t)$ for any $v \in V \setminus F$; as this also implies that $c(v, t'+1) = c(v, t')$ for all $v \in V \setminus F$ and $t' > t_0 + T(\mathbf{C})$, this will complete the proof.

Accordingly, consider any execution of Step 3ci in a round $t' > t_0 + T(\mathbf{C})$. The instance of $\mathbf{C}$ terminating in this round was started in round $t' - T(\mathbf{C}) > t_0$. However, in this round the weak pulser must have generated a pulse, yielding that, in fact, $t' - T(\mathbf{C}) \geq t_0 + T(\mathbf{C})$. Assuming for contradiction that t' is the earliest round in which the claim is violated, we thus have that $c'(v, t' - T(\mathbf{C})) = c'(w, t' - T(\mathbf{C}))$ for all $v, w \in V \setminus F$, i.e., all correct nodes used the same input value c for the instance. By the validity property of $\mathbf{C}$, this implies that $v \in V \setminus F$ outputs $y(v, t') = c$ in round t' and sets $c'(v, t') = c + T(\mathbf{C}) \bmod \Psi$. However, since t' is the earliest round of violation, we already have that $c'(v, t') = c(v, t') = c + T(\mathbf{C}) \bmod C$ after the second step, contradicting the assumption and showing that the execution stabilized in round $t_0 + T(\mathbf{C}) + 1 \leq T(\mathbf{W}) + T(\mathbf{C}) + 1$. $\qquad\square$

Together with Lemma 1, we get the following corollary.

Corollary 1. *Let $\Psi > 1$. Suppose there exists an f-resilient Ψ-value consensus routine $\mathbf{C}$ and a weak Φ-pulser $\mathbf{W}$, where $\Phi \geq T(\mathbf{C})$. Then there exists an f-resilient strong Ψ-pulser $\mathbf{P}$ that*

- *stabilises in time $T(\mathbf{P}) \leq T(\mathbf{C}) + T(\mathbf{W}) + \Psi$, and*
- *uses message of size at most $M(\mathbf{P}) \leq M(\mathbf{C}) + M(\mathbf{W})$ bits.*

5 Constructing weak pulsers from less resilient strong pulsers

Having seen that we can construct strong pulsers from weak pulsers using a consensus algorithm, the only piece missing in our framework is the existence of efficient weak pulsers. Indeed, having a pair of an f-resilient weak pulser and a consensus routine, we immediately obtain a corresponding firing squad algorithm.

In this section, we devise a recursive construction of a weak pulser from strong pulsers of smaller resilience. Given that a 0-resilient pulser is trivial and that we can obtain strong pulsers from weak ones without losing resilience, this is sufficient for constructing strong pulsers of optimal resilience from consensus algorithms of optimal resilience.

Our approach bears similarity to our constructions from earlier work [24, 26], but attains better bit complexity and can be used with an arbitrary consensus routine. On a high level, we take the following approach:

1. Partition the network into two parts, each running a strong pulser (with small resilience). Our construction guarantees that at least one of the strong pulsers stabilises.
2. Consider the pulses generated by the strong pulsers as *potential* pulses.
3. Since one of the strong pulsers may not stabilise, it may generate *spurious* pulses, that is, pulses that only a subset of the correct nodes observe.
4. We first limit the frequency of the spurious pulses using a filtering mechanism based on threshold voting.
5. We then enforce any spurious pulse to be observed by all correct nodes by employing a *silent consensus* routine. In silent consensus, no message is sent (by correct nodes) if all correct nodes have input 0. Thus, if all nodes actually participating in an instance have input 0, non-participating nodes behave *as if they participated* with input 0. This avoids the chicken-and-egg problem of having to solve consensus on participation in the consensus routine. We make sure that if any node uses input 1, i.e., the consensus routine may output 1, all nodes participate. Thus, when a pulse is generated, all correct nodes agree on this.

5.1 The filtering construction

Our goal is to construct a weak Φ-pulser (for sufficiently large Φ) with resilience f. We partition the set of n nodes into two disjoint sets V_0 and V_1 with n_0 and n_1 nodes, respectively. Thus, we have $n = n_0 + n_1$. For $i \in \{0, 1\}$, let $\mathbf{P}_i$ be an f_i-resilient strong Ψ_i-pulser. That is, $\mathbf{P}_i$ generates

a pulse every Ψ_i rounds once stabilised, granted that V_i contains at most f_i faulty nodes. Nodes in block i execute the algorithm $\mathbf{P}_i$. Our construction tolerates $f = f_0 + f_1 + 1$ faulty nodes. Since we consider Byzantine faults, we require the additional constraint that $f < n/3$.

Let $a_i(v, t) \in \{0, 1\}$ indicate the output bit of $\mathbf{P}_i$ for a node $v \in V_i$. Note that we might have a block $i \in \{0, 1\}$ that contains more than f_i faulty nodes. Thus, it is possible that the algorithm $\mathbf{P}_i$ never stabilises. In particular, we might have the situation that some of the nodes in block i produce a pulse, but others do not. We say that a pulse generated by such a $\mathbf{P}_i$ is *spurious*. We proceed by showing how to filter out such spurious pulses if they occur too often.

Filtering rules. We define five variables with the following semantics:

- $m_i(v, t + 1)$ indicates whether at least $n_i - f_i$ nodes $u \in V_i$ sent $a_i(u, t) = 1$,
- $M_i(v, t + 1)$ indicates whether at least $n - f$ nodes $u \in V$ sent $m_i(u, t) = 1$,
- $\ell_i(v, t)$ indicates when was the last time block i triggered a (possibly spurious) pulse,
- $w_i(v, t)$ indicates how long any firing events coming from block i are ignored, and
- $b_i(v, t)$ indicates whether node v accepts a firing event from block i.

The first two of the above variables are set according to the following rules:

- $m_i(v, t + 1) = 1$ if and only if $|\{u \in V_i : a_i(v, u, t) = 1|\} \geq n_i - f_i$,
- $M_i(v, t + 1) = 1$ if and only if $|\{u \in V : m_i(v, u, t) = 1\} \geq n - f$.

We update the $\ell(\cdot, \cdot)$ variables using the rule

$$\ell_i(v, t + 1) = \begin{cases} 0 & \text{if } |\{u \in V : m_i(u, t) = 1\}| \geq f + 1, \\ \min\{\Psi_i, \ell_i(v, t) + 1\} & \text{otherwise.} \end{cases}$$

In words, the counter is reset on round $t + 1$ if v has proof that at least one correct node u had $m_i(u, t) = 1$, that is, some u observed $\mathbf{P}_i$ generating a (possibly spurious) pulse.

We reset the cooldown counter w_i whenever suspicious activity occurs. The idea is that it is reset to its maximum value C by node v in the following two cases:

- some other correct node $u \neq v$ observed block i generating a pulse, but the node v did not
- block i generated a pulse, but this happened either too soon or too late.

To capture this behaviour, the cooldown counter is set with the rule

$$w_i(v, t + 1) = \begin{cases} C & \text{if } M_i(v, t + 1) = 0 \text{ and } \ell_i(v, t + 1) = 0, \\ C & \text{if } M_i(v, t + 1) = 1 \text{ and } \ell_i(v, t) \neq \Psi_i - 1, \\ \max\{w_i(v, t) - 1, 0\} & \text{otherwise,} \end{cases}$$

where $C = \max\{\Psi_0, \Psi_1\} + \Phi + 2$. Finally, a node v accepts a pulse generated by block i if the node's cooldown counter is zero and it saw at least $n - f$ nodes supporting the pulse. The variable $b_i(v, t)$ indicates whether node v accepted a pulse from block i on round t. The variable is set using the rule

$$b_i(v, t) = \begin{cases} 1 & \text{if } w_i(v, t) = 0 \text{ and } M_i(v, t) = 1, \\ 0 & \text{otherwise.} \end{cases}$$

5.2 Analysis of the filtering construction

We now analyse when the nodes accept firing events generated by the blocks. We say that a block i is correct if it contains at most f_i faulty nodes. Note that since there are at most $f = f_0 + f_1 + 1$ faulty nodes, at least one block $i \in \{0, 1\}$ will be correct. Thus, eventually the algorithm $\mathbf{P}_i$ run by a correct block i will stabilise. This yields the following lemma.

Lemma 2. *For some $i \in \{0,1\}$ the strong pulser algorithm $\mathbf{P}_i$ stabilises by round $T(\mathbf{P}_i)$.*

We proceed by establishing some bounds on when (possibly spurious) pulses generated by block i are accepted. We start with the easy case of having a correct block i.

Lemma 3. *If block i is correct, then there exists a round $t_0 \leq T(\mathbf{P}_i) + 2$ such that for each $v \in V \setminus F$, $M_i(v,t) = 1$ if and only if $t = t_0 + k\Psi_i$ for $k \in \mathbb{N}_0$.*

Proof. If block i is correct, then the algorithm $\mathbf{P}_i$ stabilises by round $T(\mathbf{P}_i)$. Hence, there is some $t_0 \leq T(\mathbf{A})$ so that

$$a(v,t) = 1 \text{ if and only if } t = t_0 + k\Psi_i \text{ for } k \in \mathbb{N}_0$$

holds for all $t \geq t_0$. If $\mathbf{P}_i$ generates a pulse on round $t \geq t_0$, then at least $n_i - f_i$ correct nodes $u \in V_i \setminus F$ have $a(u,t) = 1$. Therefore, for all $v \in V \setminus F$ we have $m_i(v,t+1) = 1$, and consequently, $M_i(v,t+2) = 1$. Since block i is correct, there are at most f_i faulty nodes in the set V_i. If $\mathbf{P}_i$ does not generate a pulse on round $t \geq t_0$, then at most $f_i < n_i - f_i$ faulty nodes[2] u may claim $a(u,t) = 1$. This yields that $m_i(v,t+1) = M_i(v,t+2) = 0$ for all $v \in V \setminus F$. $\square$

We can now establish that a correct node accepts a pulse generated by a correct block i exactly every Ψ_i rounds.

Lemma 4. *If block i is correct, then there exists a round $t_0 \leq T(\mathbf{P}_i) + 2C$ such that for each $v \in V \setminus F$, $b_i(v,t) = 1$ for any $t \geq t_0$ if and only if $t = t_0 + k\Psi_i$ for $k \in \mathbb{N}_0$.*

Proof. Lemma 3 implies that there exists $t_0' \leq T(\mathbf{P}_i) + 2$ such that both $M_i(v,t) = 1$ and $\ell_i(v,t) = 0$ hold for $t \geq t_0'$ if and only if $t = t_0' + k\Psi_i$ for $k \in \mathbb{N}_0$. Thus, it follows that $w_i(v,t+1) = \max\{w_i(v,t) - 1, 0\}$ for all such t and hence $w_i(v,t') = 0$ for all $t' \geq t_0' + C + 2$. The claim now follows from the definition of $b_i(v,t')$, the choice of t_0', and the fact that $\Psi_i \leq C - 2$. $\square$

It remains to deal with the faulty block. If we have Byzantine nodes, then a block i with more than f_i faulty nodes may attempt to generate spurious pulses. However, the filtering mechanism prevents the spurious pulses from occuring too frequently.

Lemma 5. *Let $v, v' \in V \setminus F$ and $t > 2$. Suppose $b_i(v,t) = 1$ and suppose that $t' > t$ is minimal such that $b_i(v',t') = 1$. Then $t' = t + \Psi_i$ or $t' > t + C$.*

Proof. Suppose $b_i(v,t) = 1$ for some correct node $v \in V$ and $t > 2$. Since $b_i(v,t) = 1$, $w_i(v,t) = 0$ and $M_i(v,t) = 1$. Because $M_i(v,t) = 1$, there must be at least $n - 2f > f$ correct nodes u such that $m_i(u,t-1) = 1$. Hence, $\ell_i(u,t) = 0$ for every node $u \in V \setminus F$.

Recall that $t' > t$ is minimal so that $b_i(v',t') = 1$. Again, $w_i(v',t') = 0$ and $M_i(v',t') = 1$. Moreover, since $\ell_i(v',t) = 0$, we must have $\ell_i(v',r) < \Psi_i - 1$ for all $t \leq r < t + \Psi_i - 1$. This implies that $t' \geq t + \Psi_i$, as $w_i(v',t') = 0$ and $M_i(v',t') = 1$ necessitate that $\ell_i(v',t'-1) = \Psi_i - 1$. In the event that $t' \neq t + \Psi_i$, the cooldown counter must have been reset at least once, i.e., $w_i(v',r) = C$ holds for some $t < r \leq t' - C$, implying that $t' > t + C$. $\square$

5.3 Introducing silent consensus

The above filtering mechanism prevents spurious pulses from occurring too often: if some node accepts a pulse from block i, then no node accepts a pulse from this block for at least Ψ_i rounds. We now strengthen the construction to enforce that any (possibly spurious) pulse generated by block i will be accepted by either all or no correct nodes. In order to achieve this, we employ *silent consensus*.

---

[2] Here we use that $f_i < n_i/3$ is necessary for the existence of a strong f_i-resilient pulser with n_i nodes, as a strong pulser solves consensus. This follows from Lemma 1 and the fact that counting reduces to consensus [14].

Definition 3 (Silent consensus). *We call a consensus protocol* silent*, if in each execution in which all correct nodes have input* 0*, correct nodes send no messages.*

The idea is that this enables to have consistent executions even if not all correct nodes actually take part in an execution, provided we can ensure that in this case all participating correct nodes use input 0: the non-participating nodes send no messages either, which is the exact same behavior participating nodes would exhibit. We show that silent consensus protocols can be obtained from non-silent ones using a simple transformation.

Theorem 4. *Any consensus protocol* $\mathbf{C}$ *can be transformed into a silent binary consensus protocol* $\mathbf{C}'$ *with* $T(\mathbf{C}') = T(\mathbf{C}) + 2$ *and the same resilience and message size.*

Proof. The new protocol $\mathbf{C}'$ can be seen as a "wrapper" protocol that manipulates the inputs and then lets each node decide whether it participates in an instance of the original protocol. The output of the original protocol, $\mathbf{C}$, will be taken into account only by correct nodes that participate throughout the protocol, as specified below.

In the first round of the new protocol, $\mathbf{C}'$, each participating node broadcasts its input if it is 0 and otherwise sends nothing. If a node receives fewer than $n - f$ times the value 1, it sets its input to 0. In the second round, the same pattern is applied.

Subsequently, $\mathbf{C}$ is executed by all nodes that received at least $f + 1$ messages in the first round; during the execution, any missing messages from nodes that do not participate are set to an arbitrary valid message by the receiver. If in the execution of $\mathbf{C}$ a node would have to send more bits than it would have according to the known bound $M(\mathbf{C})$, it (locally) aborts the execution of $\mathbf{C}$. Likewise, if the running time bound of $\mathbf{C}$ would be violated, it aborts as well. Finally, a node outputs 0 in the new protocol if it did not participate in the execution of $\mathbf{C}$, aborted it, or received f or fewer messages in the second round, and it outputs the result according to the run of $\mathbf{C}$ otherwise.

We first show that the new protocol, $\mathbf{C}'$, is a consensus protocol with the same resilience as $\mathbf{C}$ and the claimed bounds on communication complexity and running time. We distinguish two cases. First, suppose that all correct nodes participate in the execution of $\mathbf{C}$ at the beginning of the third round. As all nodes participate, the bounds on resilience, communication complexity, and running time that apply to $\mathbf{C}$ hold in this execution, and no node will quit executing the protocol before termination. To establish agreement and validity, again we distinguish two cases. If all nodes output the outcome of the execution of $\mathbf{C}$, these properties follow right away since $\mathbf{C}$ satisfies them; here we use that although the initial two rounds might affect the inputs of nodes, a node will change its input to 0 only if there is at least one correct node with input 0. On the other hand, if some node outputs 0 because it received f or fewer messages in the second round of $\mathbf{C}'$, no node received more than $2f < n - f$ messages in the second round. Consequently, all nodes executed $\mathbf{C}$ with input 0 and computed output 0 by the agreement property of $\mathbf{C}$, implying agreement and validity of the new protocol.

The second case is that some correct node does not participate in the execution of $\mathbf{C}$. Thus, it received at most f messages in the first round of $\mathbf{C}'$, implying that no node received more than $2f < n - f$ messages in this round. Consequently, correct nodes set their input to 0 and will not transmit in the second round. While some nodes may execute $\mathbf{C}$, all correct nodes will output 0 no matter how $\mathbf{C}$ behaves. Since nodes abort the execution of $\mathbf{C}$ if the bounds on communication or time complexity are about to be violated, the claimed bounds for the new protocol hold.

It remains to show that the new protocol is silent. Clearly, if all correct nodes have input 0, they will not transmit in the first two rounds. In particular, they will not receive more than f messages in the first round and not participate in the execution of $\mathbf{C}$. Hence correct nodes do not send messages at all, as claimed. □

For example, plugging in the phase king protocol [3], we get the following corollary.

Corollary 2. *For any $f < n/3$, there exists an f-resilient deterministic silent binary consensus protocol $\mathbf{C}$ with $T(\mathbf{C}) \in \Theta(f)$ and $M(\mathbf{C}) \in O(1)$.*

5.4 Using silent consensus to prune spurious pulses

As the filtering construction bounds the frequency at which spurious pulses may occur from below, we can make sure that at each time, only one consensus instance can be executed for each block. However, we need to further preprocess the inputs, in order to make sure that either (i) all correct nodes participate in an instance or (ii) no participating correct node has input 1; here, output 1 means agreement on a pulse being triggered, while output 0 results in no action.

Recall that $b_i(v, t) \in \{0, 1\}$ indicates whether v observed a (filtered) pulse of the strong pulser $\mathbf{P}_i$ in round t. Moreover, assume that $\mathbf{C}$ is a silent consensus protocol running in $T(\mathbf{C})$ rounds. We use two copies $\mathbf{C}_i$, where $i \in \{0, 1\}$, of the consensus routine $\mathbf{C}$. We require that $\Psi_i \geq T(\mathbf{C})$, which guarantees by Lemma 5 that (after stabilisation) every instance of $\mathbf{C}$ has sufficient time to complete. Adding one more level of voting to clean up the inputs, we arrive at the following routine.

The pruning algorithm. Besides the local variables of $\mathbf{C}_i$, the algorithm will use the following variables for each $v \in V$ and round $t \in \mathbb{N}$:

- $y_i(v, t) \in \{0, 1\}$ denotes the output value of consensus routine $\mathbf{C}_i$,
- $r_i(v, t) \in \{1, \ldots, T(\mathbf{C})\} \cup \{\bot\}$ is a local round counter for controlling $\mathbf{C}_i$, and
- $B_i(v, t) \in \{0, 1\}$ is the output of block i.

Now each node v executes the following on round t:

1. Broadcast $b_i(v, t)$.
2. If $b_i(w, t) = 1$ for at least $n - 2f$ nodes $w \in V$, then set $r_i(v, t + 1) = 1$.
3. If $r_i(v, t) = 1$, then

   (a) start a new instance of $\mathbf{C}_i$, that is, re-initialise the variables of $\mathbf{C}_i$ correctly,
   (b) use input 1 if $b_i(w, t - 1) = 1$ for at least $n - f$ nodes $w \in V$ and 0 otherwise.

4. If $r_i(v, t) = T(\mathbf{C})$, then

   (a) set $r_i(v, t + 1) = \bot$,
   (b) set $B_i(v, t + 1) = y_i(v, t)$, where $y_i(v, t) \in \{0, 1\}$ is the output variable of $\mathbf{C}_i$.

   Otherwise, set $B_i(v, t + 1) = 0$.

5. If $r_i(v, t) \notin \{T(\mathbf{C}), \bot\}$, then

   (a) execute round $r_i(v, t)$ of $\mathbf{C}_i$, and
   (b) set $r_i(v, t + 1) = r_i(v, t) + 1$.

Analysis. Besides the communication used for computing the values $b_i(\cdot)$, the above algorithm uses messages of size $M(\mathbf{C}) + 1$, as $M(\mathbf{C})$ bits are used when executing $\mathbf{C}_i$ and one bit is used to communicate the value of $b_i(v, t)$.

We say that $v \in V \setminus F$ *executes round* $r \in \{1, \ldots, T(\mathbf{C})\}$ of $\mathbf{C}_i$ in round t iff it executes the final step of the above algorithm in round t with $r_i(v, t) = r$. By Lemma 5, in rounds $t > T(\mathbf{C}) + 2$, there is always at most one instance of $\mathbf{C}_i$ being executed, and if so, consistently.

Corollary 3. *Suppose $v \in V \setminus F$ executes round 1 of $\mathbf{C}_i$ in some round $t > T(\mathbf{C}) + 2$. Then there is a subset $U \subseteq V \setminus F$ such that each $u \in U$ executes round $r \in \{1, \ldots, T(\mathbf{C})\}$ of $\mathbf{C}_i$ in round $t + r - 1$ and no $u \in V \setminus (F \cup U)$ executes any round of $\mathbf{C}_i$ in round $t + r - 1$.*

Exploiting silence of $\mathbf{C}_i$ and the choice of inputs, we can ensure that the case $U \neq V \setminus F$ causes no trouble.

Lemma 6. *Let $t > T(\mathbf{C}) + 2$ and U be as in Corollary 3. Then $U = V \setminus F$ or each $u \in U$ has input 0 for the respective instance of $\mathbf{C}_i$.*

Proof. Suppose $u \in U$ starts an instance with input 1 in round $t' \in \{t - T(\mathbf{C}) - 1, \dots, t\}$. Then $b_i(w, t' - 1) = 1$ for at least $n - 2f$ nodes $w \in V \setminus F$, since u received $b_i(w, t' - 1) = 1$ from $n - f$ nodes $w \in V$. Thus, each $v \in V \setminus F$ received $b_i(w, t' - 1) = 1$ from at least $n - 2f$ nodes w and sets $r_i(v, t') = 1$, i.e., $U = V \setminus F$. The lemma now follows from Corollary 3. $\qquad\square$

Recall that if all nodes executing $\mathbf{C}_i$ have input 0, non-participating correct nodes behave exactly as if they executed $\mathbf{C}_i$ as well, i.e., they send no messages. Hence, if $U \neq V \setminus F$, all nodes executing the algorithm will compute output 0. Therefore, Corollary 3, Lemma 5, and Lemma 6 imply the following corollary.

Corollary 4. *In rounds $t > T(\mathbf{C}) + 2$, $B_i(v, t) = B_i(w, t)$ for all $v, w \in V \setminus F$ and $i \in \{0, 1\}$. Furthermore, if $B_i(v, t) = 1$ for $v \in V \setminus F$ and $t > T(\mathbf{C}) + 2$, then the minimal $t' > t$ so that $B_i(v, t') = 1$ (if it exists) satisfies either $t' = t + \Psi_i$ or $t' > t + C = t + \max\{\Psi_0, \Psi_1\} + \Phi + 2$.*

Finally, we observe that our approach does not filter out pulses from correct blocks.

Lemma 7. *If block i is correct, there is a round $t_0 \leq T(\mathbf{P}_i) + 2C + T(\mathbf{C}) + 1$ so that for any $t \geq t_0$, $B_i(v, t) = 1$ if and only if $t = t_0 + k\Psi_i$ for some $k \in \mathbb{N}_0$.*

Proof. Lemma 4 states the same for the variables $b_i(v, t)$ and a round $t'_0 \leq T(\mathbf{P}_i) + 2C$. If $b_i(v, t) = 1$ for all $v \in V \setminus F$ and some round t, all correct nodes start executing an instance of $\mathbf{C}_i$ with input 1 in round $t + 1$. As, by Corollary 3, this instance executes correctly and, by validity of $\mathbf{C}_i$, output 1 in round $t + T(\mathbf{C})$, all correct nodes satisfy $B_i(v, t + T(\mathbf{C}) + 1) = 1$. Similarly, $B_i(v, t + T(\mathbf{C}) + 1) = 0$ for such v and any $t \geq t'_0$ with $b_i(v, t) = 0$. $\qquad\square$

5.5 Obtaining the weak pulser

Finally, we define the output variable of our weak pulser as

$$B(v, t + 1) = \max\{B_0(v, t), B_1(v, t)\}.$$

As we have eliminated the possibility that $B_i(v, t) \neq B_i(w, t)$ for $v, w \in V \setminus F$ and $t > T(\mathbf{C}) + 2$, Property W1 holds. Since there is at least one correct block i by Lemma 2, Lemma 7 shows that there will be good pulses (satisfying Properties W2 and W3) regularly, unless block $1 - i$ interferes by generating pulses violating Property W3 (i.e., in too short order after a pulse generated by block i). Here the filtering mechanism comes to the rescue: as we made sure that pulses are either generated at the chosen frequency Ψ_i or a long period of C rounds of generating no pulse is enforced (Corollary 4), it is sufficient to choose Ψ_0 and Ψ_1 as coprime multiples of Φ.

Accordingly, we pick $\Psi_0 = 2\Phi$ and $\Psi_1 = 3\Phi$ and observe that this results in a good pulse within $O(\Phi)$ rounds after the B_i stabilised.

Lemma 8. *In the construction described in the previous two subsections, choose $\Psi_0 = 2\Phi$ and $\Psi_1 = 3\Phi$ for any $\Phi \geq T(\mathbf{C})$. Then $B(v, t)$ is the output variable of a weak Φ-pulser with stabilisation time $\max\{T(\mathbf{P}_0), T(\mathbf{P}_1)\} + O(\Phi)$.*

Proof. We have that $C = \max\{\Psi_0, \Psi_1\} + \Phi + 2 \in O(\Phi)$. By the above observations, there is a round $t \in \max\{T(\mathbf{P}_0), T(\mathbf{P}_1)\} + T(\mathbf{C}) + O(\Phi) = \max\{T(\mathbf{P}_0), T(\mathbf{P}_1)\} + O(\Phi)$ satisfying the following four properties. For either block $i \in \{0, 1\}$, we have by Corollary 4 that

1. $B_i(v, t') = B_i(w, t')$ and $B(v, t') = B(w, t')$ for any $v, w \in V \setminus F$ and $t' \geq t$.

Moreover, for a correct block i and for all $v \in V \setminus F$ we have from Lemma 7 that

2. $B_i(v, t) = B_i(v, t + \Psi_i) = 1$,

18

3. $B_i(v, t') = 0$ for all $t' \in \{t + 1, \ldots, t + \Phi - 1\} \cup \{t + \Psi_i + 1, \ldots, t + \Psi_i + \Phi - 1\}$,

and for a (possibly faulty) block $1 - i$ we have from Corollary 4 that

4. if $B_{1-i}(v, t') = 1$ for some $v \in V \setminus F$ and $t' \in \{t+1, \ldots, t+\Psi_i+\Phi-1\}$, then $B_{1-i}(u, t'') = 0$ for all $u \in V \setminus F$ and $t'' \in \{t' + 1, \ldots, t' + C\}$ that do not satisfy $t'' = t' + k\Psi_{1-i}$ for some $k \in \mathbb{N}_0$.

Now it remains to argue that a good pulse is generated. Suppose that i is a correct block given by Lemma 2. By the first property, it suffices to show that a good pulse occurs in round t or in round $t + \Psi_i$. From the second property, we get for all $v \in V \setminus F$ that $B(v, t) = 1$ and $B(v, t + \Psi_i) = 1$. If the pulse in round t is good, the claim holds. Hence, assume that there is a round $t' \in \{t + 1, \ldots, t + \Psi_i - 1\}$ in which another pulse occurs, that is, $B(v, t') = 1$ for some $v \in V \setminus F$. This entails that $B_{1-i}(v, t') = 1$ by the third property. We claim that in this case the pulse in round $t + \Psi_i$ is good. To show this, we exploit the fourth property. Recall that $C > \Psi_i + \Phi$, i.e., $t' + C > t + \Psi_i + \Phi$. We distinguish two cases:

- In the case $i = 0$, we have that $t' + \Psi_{1-i} = t' + 3\Phi = t' + \Psi_0 + \Psi > t + \Psi_0 + \Phi$, that is, the pulse in round $t + \Psi_0 = t + \Psi_i$ is good.

- In the case $i = 1$, we have that $t' + \Psi_{1-i} = t' + 2\Phi < t + 3\Phi \leq t + \Psi_0$ and $t' + 2\Psi_{1-i} = t' + 4\Phi = t' + \Psi_1 + \Phi > t + \Psi_1 + \Phi$, that is, the pulse in round $t + \Psi_1 = t + \Psi_i$ is good.

In either case, a good pulse occurs by round $t + \max\{\Psi_0, \Psi_1\} \in \max\{T(\mathbf{P}_0), T(\mathbf{P}_1)\} + O(\Phi)$. $\square$

From the above lemma and the constructions discussed in this section, we get the following theorem.

Theorem 5. *Let $n = n_0 + n_1$ and $f = f_0 + f_1 + 1$, where $n > 3f$. Suppose $\mathbf{C}$ is an f-resilient consensus algorithm on n nodes and $\mathbf{P}_i$, $i \in \{0, 1\}$, are f_i-resilient strong pulser algorithms on n_i nodes. Then, for any $\Phi \geq T(\mathbf{C}) + 2$, there exists an f-resilient weak Φ-pulser $\mathbf{W}$ on n nodes that satisfies*

- $T(\mathbf{W}) \in \max\{T(\mathbf{P}_0), T(\mathbf{P}_1)\} + O(\Phi)$,
- $M(\mathbf{W}) \in \max\{M(\mathbf{P}_0), M(\mathbf{P}_1)\} + O(M(\mathbf{C}))$.

Proof. By Theorem 4, we can transform $\mathbf{C}$ into a silent consensus protocol $\mathbf{C}'$, at the cost of increasing its round complexity by 2. Using $\mathbf{C}'$ in the construction, Lemma 8 shows that we obtain a weak Φ-pulser with the stated stabilisation time, which by construction tolerates f faults. Concerning the message size, note that we run $\mathbf{P}_0$ and $\mathbf{P}_1$ on disjoint node sets. Apart from sending $\max\{M(\mathbf{P}_0), M(\mathbf{P}_1)\}$ bits per round for its respective strong pulser, each node may send $M(\mathbf{C})$ bits each to each other node for the two copies $\mathbf{C}_i$ of $\mathbf{C}$ it runs in parallel, plus a constant number of additional bits for the filtering construction including its outputs $b_i(\cdot, \cdot)$. $\square$

6 Main results

Finally, we put the developed machinery to use, recursively generating strong pulsers out of consensus algorithms. We first give deterministic algorithms and then consider randomised algorithms.

6.1 Deterministic algorithms

Theorem 6. *Suppose that we are given a family of f-resilient deterministic consensus algorithms $\mathbf{C}(f)$ running on any number $n > 3f$ of nodes in $T(\mathbf{C}(f))$ rounds using $M(\mathbf{C}(f))$-bit messages, where $T(\mathbf{C}(f))$ and $M(\mathbf{C}(f))$ are increasing in f. Then, for any $\Psi \in \mathbb{N}$, $f \in \mathbb{N}_0$, and $n > 3f$, there exists a strong Ψ-pulser $\mathbf{P}$ on n nodes that*

- *stabilises in time $T(\mathbf{P}) \in (1 + o(1))\Psi + O\left(\sum_{j=0}^{\lceil \log f \rceil} T(\mathbf{C}(2^j))\right)$*
- *uses messages of size at most $M(\mathbf{P}) \in O\left(1 + \sum_{j=0}^{\lceil \log f \rceil} M(\mathbf{C}(2^j))\right)$ bits,*

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

Proof. We show by induction on k that f-resilient strong Ψ-pulsers $\mathbf{P}(f, \Psi)$ on $n > 3f$ nodes with the stated complexity exist for any $f < 2^k$, with the addition that the (bounds on) stabilisation time and message size of our pulsers are increasing in f. We anchor the induction at $k = 0$, i.e., $f = 0$, for which, trivially, a 0-resilient strong Ψ-pulser with $n \in \mathbb{N}$ nodes is given by one node generating pulses locally and informing the other nodes when to do so. This requires 1-bit messages and stabilises in Ψ rounds.

Now assume that $2^k \le f < 2^{k+1}$ for $k \in \mathbb{N}_0$ and the claim holds for all $0 \le f < 2^k$. Since $2 \cdot (2^k - 1) + 1 = 2^{k+1} - 1$, there are $f_0, f_1 < 2^k$ such that $f = f_0 + f_1 + 1$. Since $n > 3f > 3f_0 + 3f_1$, we can pick $n_i > 3f_i$ for both $i \in \{0, 1\}$ satisfying $n = n_0 + n_1$. Denoting by $\mathbf{P}(f', \Psi')$ the strong Ψ'-pulsers that exist by the induction hypothesis for $f' < 2^k$, we thus can apply Theorem 5 to $\mathbf{C}(f)$, $\mathbf{P}_i = \mathbf{P}(f_i, \Phi)$, and $\Phi \in O(\log \Psi) + T(\mathbf{C}(f))$ chosen in accordance with Theorem 1 for $L = \Psi$. Note that without loss of generality, the $O(\log \Psi)$ term is at least 2, that is, $\Phi \ge T(\mathbf{C}(f)) + 2$ holds as required.

We obtain a weak Φ-pulser $\mathbf{W}$ with resilience f on n nodes and

- $T(\mathbf{W}) \in \max\{T(\mathbf{P}(f_0, \Phi)), T(\mathbf{P}(f_1, \Phi))\} + O(\Phi) \subseteq T(\mathbf{P}(2^k, \Phi)) + O(\log \Psi + T(\mathbf{C}(f)))$,
- $M(\mathbf{W}) \in \max\{M(\mathbf{P}(f_0, \Phi)), M(\mathbf{P}(f_1, \Phi))\} + O(M(\mathbf{C}(f))) \subseteq M(\mathbf{P}(2^k, \Phi)) + O(M(\mathbf{C}(f)))$.

Next, we apply Theorem 1 to $\mathbf{C}(f)$ to obtain an f-resilient Ψ-value consensus protocol that uses $M(\mathbf{C}(f))$-bit messages and runs in Φ rounds. We feed the weak pulser and this consensus protocol into Corollary 1 to obtain an f-resilient strong Ψ-pulser $\mathbf{P}$ with

- $T(\mathbf{P}) \le T(\mathbf{W}) + T(\mathbf{C}(f)) + \Psi \in T(\mathbf{P}(2^k, \Phi)) + (1 + o(1))\Psi + O(T(\mathbf{C}(f)))$,
- $M(\mathbf{P}) \le M(\mathbf{W}) + M(\mathbf{C}(f)) \in M(\mathbf{P}(2^k, \Phi)) + O(M(\mathbf{C}(f)))$.

Plugging in the bounds given by the induction hypothesis, and using that we get that $T(\mathbf{C}(f))$ and $M(\mathbf{C}(f))$ are increasing in f, we get that

- $T(\mathbf{W}) \in (1 + o(1))\Psi + O\left(\sum_{j=0}^{\lceil \log f \rceil} T(\mathbf{C}(2^j))\right)$,
- $M(\mathbf{W}) \in O\left(1 + \sum_{j=0}^{\lceil \log f \rceil} M(\mathbf{C}(2^j))\right)$.

Because we bounded complexities using $T(\mathbf{P}(2^k, \Phi))$ and $M(\mathbf{P}(2^k, \Phi))$, and $T(\mathbf{C}(f))$ and $M(\mathbf{C}(f))$ are increasing in f by assumption, we also maintain that the new bounds are increasing in f. Thus, the induction step succeeds and the proof is complete. $\square$

Plugging in the phase king protocol [3], which has optimal resilience, running time $O(f)$, and constant message size, we can extract a strong pulser that is optimally resilient, has asymptotically optimal stabilization time, and message size $O(\log f)$.

Corollary 5. *For any $\Psi, f \in \mathbb{N}$ and $n > 3f$, an f-resilient strong Ψ-pulser on n nodes with stabilisation time $(1 + o(1))\Psi + O(f)$ and message size $O(\log f)$ exists.*

We obtain efficient solutions to the firing squad and synchronous counting problems.

Corollary 6. *For any $f \in \mathbb{N}$ and $n > 3f$, an f-resilient firing squad on n nodes with stabilisation and response times of $O(f)$ and message size $O(\log f)$ exists.*

Proof. We use Corollary 5 with $\Psi \in O(f)$ being the running time of the phase king protocol [3], followed by applying Theorem 2 to the obtained pulser and the phase king protocol. $\square$

Corollary 7. *For any $C, f \in \mathbb{N}$ and $n > 3f$, an f-resilient C-counter on n nodes with stabilisation time $O(f + \log C)$ and message size $O(\log f)$ exists.*

Proof. In the last step of the construction of Theorem 6, we do not use Corollary 1 to extract a strong pulser, but directly obtain a counter using Theorem 3. This avoids the overhead of Ψ due to waiting for the next pulse. Recalling that the $o(\Psi)$ term in the complexity comes from the $O(\log \Psi)$ additive overhead in time of the multi-value consensus routine, the claim follows. $\square$

We remark that one can strengthen the bound on the stabilisation time to $O(f + (\log \Psi)/B)$ using messages of size B, by using larger messages in the reduction given by Theorem 1 [25]. However, this affects the asymtotic stabilisation time only if Ψ is super-exponential in f.

6.2 Randomisation

While we confined our discussion to the deterministic setting, it is straightforward to adapt our solution to utilising randomised consensus routines. For our framework, we require that the running time of the algorithms satisfy deterministic bounds, while we allow for a probabilistic guarantee on the agreement and validity properties. More precisely, we demand that these properties hold with some probability $1 - p$, where $p \leq 1/f^c$ for a sufficiently large constant c. Noting that our recursive construction of strong pulsers involves $f^{O(1)}$ calls to the utilised consensus routine within $f^{O(1)}$ rounds, a union bound shows that with probability $1 - 1/f^{c-O(1)}$ in all instances agreement and validity hold. This yields the following generalisation of Theorem 6.

Theorem 7. *Suppose that for constant $\varepsilon \geq 0$ we are given a family of f-resilient consensus algorithms $\mathbf{C}(f)$ running on any number $n > (3 + \varepsilon)f$ of nodes in $T(\mathbf{C}(f))$ rounds using $M(\mathbf{C}(f))$-bit messages, where $T(\mathbf{C}(f))$ and $M(\mathbf{C}(f))$ are increasing in f, and $\mathbf{C}(f)$ succeeds with probability $p \geq 1 - 1/n^c$ for sufficiently large $c \in O(1)$. Then, for any $\Psi, f \in \mathbb{N}$ and $n > (3 + \varepsilon)f$, a strong probabilistic Ψ-pulser $\mathbf{S}$ on n nodes with*

- $T(\mathbf{S}) \in (1 + o(1))\Psi + O\left(\sum_{j=0}^{\lceil \log f \rceil} T(\mathbf{C}(2^j))\right)$
- $M(\mathbf{S}) \in O\left(1 + \sum_{j=0}^{\lceil \log f \rceil} M(\mathbf{C}(2^j))\right)$

exists, where for $f = 0$ the sums are empty. Here, probabilistic means that the pulser may fail to behave correctly in round $t \geq T(\mathbf{S})$ with probability $f^{O(1)}p$ (and then needs to re-stabilise).

The additional reservation that $\mathbf{C}$ may require $n > (3 + \varepsilon)f$ accounts for the fact that various randomised consensus protocols have slightly suboptimal resilience. Note also that randomised protocols may have further requirements on the system, like private channels or a static adversary (faulty nodes are determined before randomness is fixed). As an example, we plug in the algorithm by King and Saia [22].

Corollary 8. *Assume that the nodes have unique identifiers such that a receiver can identify the sender of a message. Moreover, suppose that communication is via private channels, i.e., faulty nodes behavior in round t is a function of all communication from correct nodes to faulty nodes in rounds $t' \leq t$. Then, for any $\Psi, f \in \mathbb{N}$ and $n > (3 + \varepsilon)f$, an f-resilient strong probabilistic Ψ-pulser on n nodes with stabilisation time $(1 + o(1))\Psi + \text{polylog } n$ and message size $\text{polylog } n$ exists, where $p = n^{-c}$ for any fixed choice of $c \in O(1)$.*

We remark that the algorithm from [22] actually limits the number of bits sent by each node to $O(\sqrt{n} \, \text{polylog } n)$, but in our framework each node broadcasts $\Omega(\log f)$ bits per round.

6.3 Crash and omission faults

Our framework can be adapted for e.g. crash and omission faults, which allow for higher resilience. Apart from replacing the utilised consensus routines, however, we need to adapt the various voting steps to the weaker fault models as well. This is a straightforward technical exercise, which is deferred to Appendix A. We confine ourselves to stating the results for deterministic firing squads here and again use the phase king protocol which tolerates $f < n/2$ omission faults and $f < n$ crash faults while having asymptotically optimal stabilisation time and message size.

Corollary 9. *For any $f \in \mathbb{N}$ and $n > 2f$, there exists an f-omission-resilient firing squad on n nodes with stabilisation and response times of $O(f)$ and message size $O(\log f)$.*

Corollary 10. *For any $f \in \mathbb{N}$ and $n > f$, there exists an f-crash-resilient firing squad on n nodes with stabilisation and response times of $O(f)$ and message size $O(\log f)$.*

Acknowledgements

References

[1] Dana Angluin, Michael J. Fischer, and Hong Jiang. Stabilizing consensus in mobile networks. In *Proc. 2nd IEEE International Conference on Distributed Computing in Sensor Systems*, volume 4026 LNCS, pages 37–50, 2006. `doi:10.1007/11776178_3`.

[2] Michael Ben-Or, Danny Dolev, and Ezra N. Hoch. Fast self-stabilizing Byzantine tolerant digital clock synchronization. In *Proc. 27th Annual ACM Symposium on Principles of Distributed Computing (PODC 2008)*, pages 385–394. ACM Press, 2008. `doi:10.1145/1400751.1400802`.

[3] Piotr Berman, Juan A. Garay, and Kenneth J. Perry. Towards optimal distributed consensus. In *Proc. 30th Annual Symposium on Foundations of Computer Science (FOCS 1989)*, pages 410–415. IEEE, 1989. `doi:10.1109/SFCS.1989.63511`.

[4] Roderick Bloem, Nicolas Braud-Santoni, and Swen Jacobs. Synthesis of self-stabilising and Byzantine-resilient distributed systems. In *Proc. 28th International Conference on Computer Aided Verification (CAV 2016), Part I*, pages 157–176, 2016. `doi:10.1007/978-3-319-41528-4_9`.

[5] James E. Burns and Nany A. Lynch. The Byzantine firing squad problem. *Advances in Computing Research*, 4:147–161, 1987.

[6] B.A. Coan, Danny Dolev, Cynthia Dwork, and Larry Stockmeyer. The distributed firing squad problem. *SIAM Journal on Computing*, 18(5):990–1012, 1989. `doi:10.1137/0218068`.

[7] Brian A. Coan and Cynthia Dwork. Simultaneity is harder than agreement. *Information and Computation*, 91(2):205–231, 1991. `doi:10.1016/0890-5401(91)90067-C`.

[8] Ariel Daliot and Danny Dolev. Self-stabilizing Byzantine agreement. In *Proc. 25th Annual ACM Symposium on Principles of Distributed Computing (PODC 2006)*, pages 143–152. ACM, 2006. `doi:10.1145/1146381.1146405`.

[9] Benjamin Doerr, Leslie Ann Goldberg, Lorenz Minder, Thomas Sauerwald, and Christian Scheideler. Stabilizing consensus with the power of two choices. In *Proc. 23rd ACM Sypomsoium on Parallelism in Algorithms and Architectures (SPAA 2011)*, pages 149–158. ACM, 2011. `doi:10.1007/978-3-642-15763-9_50`.

[10] Danny Dolev. The Byzantine generals strike again. *Journal of Algorithms*, 3(1):14–30, 1982.

[11] Danny Dolev and Ezra N. Hoch. On self-stabilizing synchronous actions despite Byzantine attacks. In *Proc. 21st International Symposium on Distributed Computing (DISC 2007)*, volume 4731 of *Lecture Notes in Computer Science*, pages 193–207. Springer, 2007. `doi: 10.1007/978-3-540-75142-7_17`.

[12] Danny Dolev and Rüdiger Reischuk. Bounds on information exchange for Byzantine agreement. *Journal of the ACM*, 32(1):191–204, 1985. `doi:10.1145/2455.214112`.

[13] Danny Dolev, Ezra N. Hoch, and Yoram Moses. An optimal self-stabilizing firing squad. *SIAM Journal on Computing*, 41(2):415–435, 2012. `doi:10.1137/090776512`.

[14] Danny Dolev, Matthias Függer, Christoph Lenzen, Ulrich Schmid, and Andreas Steininger. Fault-tolerant distributed systems in hardware. *Bulletin of the EATCS*, (116), 2015. `http://bulletin.eatcs.org/index.php/beatcs/issue/view/18`.

[15] Danny Dolev, Keijo Heljanko, Matti Järvisalo, Janne H. Korhonen, Christoph Lenzen, Joel Rybicki, Jukka Suomela, and Siert Wieringa. Synchronous counting and computational algorithm design. *Journal of Computer and System Sciences*, 82(2):310–332, 2016. `doi: 10.1016/j.jcss.2015.09.002`.

[16] Shlomi Dolev. *Self-Stabilization*. The MIT Press, Cambridge, MA, 2000.

[17] Shlomi Dolev and Jennifer L. Welch. Self-stabilizing clock synchronization in the presence of Byzantine faults. *Journal of the ACM*, 51(5):780–799, 2004. `doi:10.1145/1017460.1017463`.

[18] Shlomi Dolev, Ronen I. Kat, and Elad M. Schiller. When consensus meets self-stabilization. *Journal of Computer and System Sciences*, 76(8):884–900, 2010. `doi:10.1016/j.jcss.2010.05.005`.

[19] Cynthia Dwork and Yoram Moses. Knowledge and common knowledge in a Byzantine environment: Crash failures. *Information and Computation*, 88(2):156–186, 1990. `doi:10.1016/0890-5401(90)90014-9`.

[20] Michael J. Fischer and Nancy A. Lynch. A lower bound for the time to assure interactive consistency. *Information Processing Letters*, 14(4):183–186, 1982. `doi:10.1016/0020-0190(82)90033-3`.

[21] Ezra N. Hoch, Danny Dolev, and Ariel Daliot. Self-stabilizing Byzantine digital clock synchronization. In *Proc. 8th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS 2006)*, volume 4280 of *Lecture Notes in Computer Science*, pages 350–362. Springer, 2006. `doi:10.1007/978-3-540-49823-0_25`.

[22] Valerie King and Jared Saia. Breaking the $O(n^2)$ bit barrier. *Journal of the ACM*, 58(4):1–24, 2011. `doi:10.1145/1989727.1989732`.

[23] Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982. `doi:10.1145/357172.357176`.

[24] Christoph Lenzen and Joel Rybicki. Efficient counting with optimal resilience. In *Proc. 29th International Symposium on Distributed Computing (DISC 2015), Tokyo, Japan, October 7–9, 2015*, volume 9363 of *Lecture Notes in Computer Science*, pages 16–30. Springer, 2015. `doi:10.1007/978-3-662-48653-5_2`.

[25] Christoph Lenzen, Matthias Függer, Markus Hofstätter, and Ulrich Schmid. Efficient construction of global time in SoCs despite arbitrary faults. In *Proc. 16th Euromicro Conference on Digital System Design (DSD 2013)*, pages 142–151, 2013. `doi:10.1109/DSD.2013.97`.

[26] Christoph Lenzen, Joel Rybicki, and Jukka Suomela. Towards optimal synchronous counting. In *Proc. 34th ACM Symposium on Principles of Distributed Computing (PODC 2015)*, pages 441–450. ACM, 2015.

[27] Christoph Lenzen, Joel Rybicki, and Jukka Suomela. Efficient counting with optimal resilience, 2016. Manuscript (extended and revised version of [24, 26]). `arXiv:1508.02535`.

[28] Nancy A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers, San Francisco, 1996.

[29] Yasuaki Nishitani and Namio Honda. The firing squad synchronization problem for graphs. *Theoretical Computer Science*, 14(1):39–61, 1981. `doi:10.1016/0304-3975(81)90004-9`.

[30] Marshall C. Pease, Robert E. Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *Journal of the ACM*, 27(2):228–234, 1980. `doi:10.1145/322186.322188`.

[31] Michel Raynal. *Fault-tolerant agreement in synchronous message-passing systems*. Morgan & Claypool, 2010. `doi:10.2200/S00294ED1V01Y201009DCT003`.

A Weaker fault types

In this section, we discuss how to adapt our framework to restricted types of permanent faults. We confine our attention to two types frequently considered in the literature, crash and omission faults [31]. We start the discussion with crash faults.

A.1 Crash faults

We now define the synchronous counting and firing squad problems under crash faults.

Definition 4 (Crash faults). *A crashing node stops executing the algorithm in some round $r \in \mathbb{N}$. In this round, the node manages to send only a subset of the messages it would send if it ran correctly. Thus, only a subset of the respective recipients receive a message from the crashed node in this round. The remaining nodes (and, in rounds $r' > r$ all nodes) receive no message.*

Crash faults are the most benign type of faults typically considered; the node crashes cleanly and causes no further trouble. This enables to use more strict requirements in the synchronous counting and firing squad problems.

In the following, let us use $F(t') \subseteq V$ to denote the set of nodes that have crashed before or in round t'. In the case of synchronous C-counting, the problem is defined similarly as in the case of Byzantine faults, but with the requirement that agreement and consistency are satisfied by the set of currently non-crashed nodes.

Definition 5 (Counting with crash faults). *In synchronous C-counting with crash faults, an execution of an algorithm stabilises in round $t \in \mathbb{N}$ if and only if all $t \leq t' \in \mathbb{N}$ the output counters $c(\cdot)$ satisfy*

- **Agreement:** $c(v, t') = c(w, t')$ *for all* $v, w \in V \setminus F(t')$ *and*
- **Consistency:** $c(v, t' + 1) = c(v, t') + 1 \bmod C$ *for all* $v, w \in V \setminus F(t' + 1)$.

For the firing squad problem, we adapt the agreement, safety, and liveness condition similarly.

Definition 6 (Firing squad with crash faults). *In the firing squad problem with crash faults, we say that an execution of an algorithm stabilises in round $t \in \mathbb{N}$ if the following three properties hold:*

- **Agreement:** $\mathrm{FIRE}(v, t') = \mathrm{FIRE}(w, t')$ *for all* $v, w \in V \setminus F(t')$ *and* $t \leq t' \in \mathbb{N}$.
- **Safety:** *If* $\mathrm{FIRE}(v, t_F) = 1$ *for* $v \in V \setminus F(t_F)$ *and* $t \leq t_F \in \mathbb{N}$, *then there is* $t_F \geq t_G \in \mathbb{N}$ *such that*

  (i) $\mathrm{GO}(w, t_G) = 1$ *for some* $w \in V \setminus F(t_G)$, *and*

  (ii) $\mathrm{FIRE}(v, t') = 0$ *for all* $t' \in \{t_G + 1, \ldots, t_F - 1\}$.

- **Liveness:** *If* $\mathrm{GO}(v, t_G) = 1$ *for* $v \in V \setminus F(t_G + 1)$ *and* $t \leq t_G \in \mathbb{N}$, *then* $\mathrm{FIRE}(v, t_F) = 1$ *for all nodes* $v \in V \setminus F(t_F)$ *and some* $t_G < t_F \in \mathbb{N}$.

A.2 Omission faults

We now turn our attention to omission faults. The case of omission faults is more challenging than crash faults, as faulty nodes may continue to participate in the execution of the algorithm. For simplicity, we focus here only on send omission faults, as our primary goal here is demonstrating the flexibility our approach. One could also consider also receive or general omission faults [31].

Definition 7 (Omission faults). *We say that a node $v \in V$ suffers from (send) omission faults if in each round r the messages sent by v are only received by some (arbitary) subset $U(r) \subseteq V$ of the nodes only. The remaining nodes in $V \setminus U(r)$ receive no message from v.*

Note that under send omission faults, the faulty nodes still receive messages from correct nodes. Hence, we modify the synchronous counting and firing squad problems as follows.

Definition 8 (Counting with omission faults)**.** *In the synchronous C-counting problem with omission faults, we require that the agreement and consistency conditions are satisfied by all nodes.*

Definition 9 (Firing squad with omission faults)**.** *In the firing squad problem with omission faults, the agreement, safety, and liveness conditions are adapted as follows. We say that an execution of an algorithm stabilises in round $t \in \mathbb{N}$ if the following three properties hold:*

- **Agreement:** $\mathrm{FIRE}(v, t') = \mathrm{FIRE}(w, t')$ *for all* $v, w \in V$ *and* $t \leq t' \in \mathbb{N}$.
- **Safety:** *If* $\mathrm{FIRE}(v, t_F) = 1$ *for* $v \in V$ *and* $t \leq t_F \in \mathbb{N}$, *then there is* $t_F \geq t_G \in \mathbb{N}$ *such that*

  (i) $\mathrm{GO}(w, t_G) = 1$ *for some* $w \in V$,

  (ii) $\mathrm{FIRE}(v, t') = 0$ *for all* $t' \in \{t_G + 1, \ldots, t_F - 1\}$.

- **Liveness:** *If* $\mathrm{GO}(v, t_G) = 1$ *for* $v \in V \setminus F$ *and* $t \leq t_G \in \mathbb{N}$, *then* $\mathrm{FIRE}(v, t_F) = 1$ *for all nodes* $v \in V$ *and some* $t_G < t_F \in \mathbb{N}$.

Finally, we remark that also the definition of consensus needs to adapted in the case of crash and send omission faults. For send omission faults, termination, agreement, and validity apply to all nodes. In the case of crash faults, these conditions need to be satisfied by all non-crashed nodes, respectively.

A.3 Necessary adjustments to the framework

While our framework could be used as is with minimal modifications, the weaker fault types permit to tolerate a larger number of faults: any $f < n/2$ in case of omission faults, and $f < n$ for crash faults. To this end, we can readily employ consensus protocols tailored to these fault types from the literature. However, we also need to adapt the voting schemes used in our constructions outside the consensus routines. Moreover, we must adapt our reductions of multivalue consensus and silent consensus to standard binary consensus. More precisely, for each fault type, we need to address and handle the following issues:

1. We need an equivalent of Theorem 1.
2. Distributing knowledge of GO inputs of 1 in the firing squad algorithm.
3. Adapting the filtering construction.
4. We need an equivalent of Theorem 4.
5. Adapting the pruning algorithm.

We now discuss each these steps for crash and omission faults separately.

A.4 Handling crash faults

For crash faults, we have to make the following modifications.

1. Each node broadcasts its input bit by bit. Denote by $U_v \subseteq V$ the set of nodes $v \in V$ did not see crashing during this phase of the algorithm. For each $u \in U_v$, v knows its input x_u. Let each node broadcast for one additional round, say a 0, and denote by $U'_v \subseteq U_v$ the set of nodes v did not observe to crash. Run binary consensus, where $v \in V$ uses input 1 if $x_u = x_{u'}$ for all $u, u' \in U_v$ and 0 else. If the output is 1, there was some node w using input 1, implying that $U_w \supseteq U'_v$ satisfies that $x_u = x_{u'}$ for all $u, u' \in U_w$; note that any non-crashed node is in $U'_v \subseteq U_w$. If the output is 1, it thus satisfies agreement to output x_v. Otherwise, output 0, trivially ensuring agreement. Concerning validity, observe that if all nodes have the same input, all nodes will use input 1 for the binary consensus routine.

2. In the absence of Byzantine faults, we may simply replace the threshold of receiving $GO(w, t-1) = 1$ from $f+1$ nodes $w \in V$ with a new the threshold of receiving $GO(w, t-1) = 1$ from any node $w \in V$ in Step 2 of the firing squad algorithm. The adjustments to the proof of Theorem 2 are straightforward.

3. We drop the requirement that $f < n/3$ from the construction. We change the rules for setting m_i, M_i, and ℓ_i as follows:

   - $m_i(v, t+1) = 1$ if no $u \in V_i$ sent $a_i(u, t) = 0$,
   - $M_i(v, t+1) = 1$ if no $u \in V$ sent $m_i(u, t) = 0$, and
   - $\ell_i(v, t+1)$ is set to 0 if there is any node u sending $m_i(u, t) = 1$.

   One can readily check that this does not affect the correctness of Lemma 2, Lemma 3, or Lemma 4. Lemma 5 is not required, as any number of faults can be tolerated; if all nodes in a block crash, clearly this block will not interfere with stabilisation, as the block will not trigger further pulses.

4. We follow the same strategy as for the Byzantine case. However, in the first two rounds, a node sets its input to 0 only if receiving no message at all. Any node receiving a message in the first round participates in the execution of the (non-silent) binary consensus protocol. Each node returns 0 if it received no message in the second round or the binary consensus protocol returned 0. If a node does not participate, all participating nodes use input 0. Note that the resulting execution is equivalent to one where all nodes participate, but the non-partipating nodes crash immediately. Hence, the participating nodes output 0 and return 0. If all nodes participate, agreement follows immediately. Concerning validity and silence, in case all nodes have input 0 no messages are sent and all nodes output 0; if all inputs are 1, no node changes its input, all nodes participate, and validity of the utilised binary consensus protocol ensures that all nodes output 1.

5. We modify Step 1 of the pruning algorithm to set $r_i(v, t+1) = 1$ if received $b_i(w, t) = 1$ from any $w \in V$, and Step 2(b) to use input 1 if not receiving $b_i(w, t-1) = 0$ from any $w \in V$. It follows that if any node $v \in V$ uses input 1 for a consensus instance whose first round is simulated in round t, each non-crashed node received $b_i(v, t-1) = 1$ and thus participates in the instance. Moreover, if all nodes have $b_i(v, t-1) = 1$, all use input 1 for the instance. Similar reasoning to the Byzantine case now establishes the required properties of the pruning routine.

None of the above modifications change message size or time bounds, implying that we can feed the modified machinery with an arbitrary binary consensus algorithm resilient to $f < n$ crash faults to obtain analogous results.

Theorem 8. *Suppose that we are given a family of f-crash-resilient deterministic consensus algorithms $\mathbf{C}(f)$ running on any number $n > f$ of nodes in $T(\mathbf{C}(f))$ rounds using $M(\mathbf{C}(f))$-bit messages, where $T(\mathbf{C}(f))$ and $M(\mathbf{C}(f))$ are increasing in f. Then, for any $\Psi, f \in \mathbb{N}$ and $n > f$, a strong Ψ-pulser $\mathbf{P}$ on n nodes with*

- $T(\mathbf{P}) \in (1 + o(1))\Psi + O\left(\sum_{j=0}^{\lceil \log f \rceil} T(\mathbf{C}(2^j))\right)$
- $M(\mathbf{P}) \in O\left(1 + \sum_{j=0}^{\lceil \log f \rceil} M(\mathbf{C}(2^j))\right)$

exists, where for $f = 0$ the sums are empty.

Plugging in a folklore algorithm for $f < n$ faults (the first crash-resilient algorithm described in [31]), we obtain the following results.

Corollary 11. *For any $\Psi, f \in \mathbb{N}$ and $n > f$, there exists a deterministic f-crash-resilient strong Ψ-pulser on n nodes with stabilisation time $(1 + o(1))\Psi + O(f)$ and message size of $O(\log f)$ bits.*

Corollary 12. *For any $f \in \mathbb{N}$ and $n > f$, there exists a deterministic f-crash-resilient firing squad on n nodes with stabilisation and response times of $O(f)$ and message size of $O(\log f)$ bits.*

Corollary 13. *For any $C, f \in \mathbb{N}$ and $n > 3f$, there exists a deterministic f-resilient C-counter on n nodes with stabilisation time $O(f + \log \Psi)$ and message size of $O(\log f)$ bits.*

A.5 Handling omission faults

We require that $n > 2f$ in the following, which is necessary and sufficient to achieve consensus in the presence of omission faults [31]. We need to take the following steps.

1. Each node broadcasts its input bit by bit. There is a unique input x that can be received $n - f > f$ times by any node (the threshold must be met for every bit, but the senders may differ). If v receives such an input, it stores it and sends it again bit-by-bit; if not, it sends nothing in this second transmission. If x is receveived at least $n - f$ times by $v \in V$ in this second iteration, v uses input 1 in a call to the binary consensus routine, otherwise 0. If v received *any* value x in this second transmission, it returns it in case the consensus routine outputs 1. If the routine outputs 0, it returns 0. Note that if any node used input 1, it received x $f + 1$ times in the second iteration, entailing that every node received x. Thus, agreement holds by the properties of the binary consensus routine. Likewise, validity of the latter implies validity of the former: if all nodes have the same input x, it is received $n - f$ times by each node in both iterations.

2. Again, we replace the threshold of receiving $\text{GO}(w, t - 1) = 1$ from $f + 1$ nodes $w \in V$ with the threshold of receiving $\text{GO}(w, t - 1) = 1$ from any node $w \in V$ in Step 2 of the firing squad algorithm, and adjusting the proof of Theorem 2 is straightforward.

3. We replace the requirement that $f < n/3$ by $f < n/2$. The only change is that $\ell_i(v, t + 1)$ is set to 0 if there is any node u sending $m_i(u, t) = 1$. One can readily check that this does not affect the correctness of Lemma 2, Lemma 3, or Lemma 4. Concerning Lemma 5, observe that any node having $b_i(v, t) = 1$ implies $M_i(v, t) = 1$ and thus $m_i(w, t - 1) = 1$ for at least $n - f > f$ nodes $w \in V$. Hence, each node $u \in V$ receives $m_i(w, t - 1) = 1$ from at least one node $w \in V$ and sets $l_i(u, t) = 0$. Lemma 5 now follows by similar reasoning as in the Byzantine case.

4. We follow the same strategy as for the Byzantine case. In the first two rounds, a node sets its input to 0 if receiving fewer than $n - f$ times 1. Any node receiving a message in the first round participates in the execution of the (non-silent) binary consensus protocol. Each node returns 0 if it received no message in the second round, it was forced to abort the binary consensus protocol due to violation of message size bound or an otherwise invalid execution, or the binary consensus protocol returned 0. If a node does not participate, there are at most f nodes with non-zero input, implying that no node receives a message in the second round. Thus, agreement holds in this case. If a node uses input one for the call to the non-silent consensus routine, all nodes participate, as at least $f + 1$ nodes sent 1 in the first round. Thus agreement follows from the correct execution of the non-silent protocol. Silence and validity are easily verified.

5. We modify Step 1 of the pruning algorithm to set $r_i(v, t + 1) = 1$ if received $b_i(w, t) = 1$ from any $w \in V$. It follows that if any node $v \in V$ uses input 1 for a consensus instance whose first round is simulated in round t, each node received $b_i(v, t - 1) = 1$ and thus participates in the instance. Moreover, if all nodes have $b_i(v, t - 1) = 1$, all use input 1

for the instance. Similar reasoning to the Byzantine case now establishes the required properties of the pruning routine.

None of the above modifications change message size or time bounds, implying that we can feed the modified machinery with an arbitrary binary consensus algorithm resilient to $f < n/2$ omission faults to obtain analogous results.

Theorem 9. *Suppose that we are given a family of f-omission-resilient deterministic consensus algorithms $\mathbf{C}(f)$ running on any number $n > 2f$ of nodes in $T(\mathbf{C}(f))$ rounds using $M(\mathbf{C}(f))$-bit messages, where $T(\mathbf{C}(f))$ and $M(\mathbf{C}(f))$ are increasing in f. Then, for any $\Psi, f \in \mathbb{N}$ and $n > 2f$, a strong f-omission-resilient Ψ-pulser $\mathbf{P}$ on n nodes with*

- $T(\mathbf{P}) \in (1 + o(1))\Psi + O\left(\sum_{j=0}^{\lceil \log f \rceil} T(\mathbf{C}(2^j))\right)$
- $M(\mathbf{P}) \in O\left(1 + \sum_{j=0}^{\lceil \log f \rceil} M(\mathbf{C}(2^j))\right)$

exists, where for $f = 0$ the sums are empty.

Plugging in a folklore algorithm for $f < n/2$ faults (the first omission-resilient algorithm described in [31]), we obtain the following results.

Corollary 14. *For any $\Psi, f \in \mathbb{N}$ and $n > 2f$, there exists a deterministic f-omission-resilient strong Ψ-pulser on n nodes with stabilisation time $(1 + o(1))\Psi + O(f)$ and message size of $O(\log f)$ bits.*

Corollary 15. *For any $f \in \mathbb{N}$ and $n > f$, there exists a deterministic f-crash-resilient firing squad on n nodes with stabilisation and response times of $O(f)$ and message size of $O(\log f)$ bits.*

Corollary 16. *For any $C, f \in \mathbb{N}$ and $n > 3f$, there exists a deterministic f-resilient C-counter on n nodes with stabilisation time $O(f + \log \Psi)$ and message size of $O(\log f)$ bits.*