

Fire!*

Krisztina Fruzsa[†]

Roman Kuznets

Ulrich Schmid

TU Wien

{kfruzsa,rkuznets,s}@ecs.tuwien.ac.at

In this paper, we provide an epistemic analysis of a simple variant of the fundamental consistent broadcasting primitive for byzantine fault-tolerant asynchronous distributed systems. Our Firing Rebels with Relay (FRR) primitive enables agents with a local preference for acting/not acting to trigger an action (FIRE) at all correct agents, in an all-or-nothing fashion. By using the epistemic reasoning framework for byzantine multi-agent systems introduced in our TARK’19 paper, we develop the necessary and sufficient state of knowledge that needs to be acquired by the agents in order to FIRE. It involves eventual common hope (a modality related to belief), which we show to be attained already by achieving eventual mutual hope in the case of FRR. We also identify subtle variations of the necessary and sufficient state of knowledge for FRR for different assumptions on the local preferences.

1 Motivation and Background

In their PODC’18 paper “Silence” [13], Goren and Moses introduced and epistemically analyzed *silent choirs* as a fundamental primitive for message-optimal protocols in synchronous fault-tolerant distributed systems where computing nodes (agents¹) can crash. In synchronous systems, where one can time-out messages, it is well-known [20] that an agent can convey information also by *not* sending some message. In a system where the sender may also crash, however, the receiver cannot infer this information from not receiving the message. Still, if only up to f of the $n > f$ agents in a system may crash, a silent choir of $f + 1$ agents that convey identical information suffices: at least one agent in the choir must be correct, so its silence can be relied on.

Whereas silent choirs also work in systems where the faulty agents may behave arbitrarily (byzantine [21]), the problem of not conveying information faithfully now also plagues messages that *are* sent, as they could originate from a faulty sender or forwarding agent. In this paper, we will introduce and epistemically analyze a fundamental primitive *Firing Rebels with Relay* (FRR), which nicely captures exactly these issues. It is a simplified version of the *consistent broadcasting* primitive introduced by Srikanth and Toueg in [27], which has been used as a pivotal building block in distributed algorithms for byzantine fault-tolerant clock synchronization [6, 11, 26, 27, 30] and synchronous consensus [28], for example.

Informally, FRR requires that *every* correct agent perform an action called FIRE, in an all-or-none fashion (though not necessarily simultaneously), and only if at least one correct agent locally observed a trigger event called START. Note that we have replaced the need to broadcast explicit information by just triggering an action, which makes FRR essentially a non-synchronous variant of the Firing Squad problem [4], hence its name. In crash-prone systems, FRR is trivial to solve, even for large f : Indeed, every agent who observes START or receives a notification message (for the first time) just invokes FIRE and sends a notification message to everyone. This guarantees that if a single correct agent observes START,

*Funded by the Austrian Science Fund (FWF) project ByzDEL P33600.

[†]PhD student in the FWF doctoral program LogiCS (W1255).

¹Since distributed systems are just one instance of multi-agent systems, we will use the term “agent” instead of “process.”

every correct agent will invoke FIRE (agents that crash during the run may or may not issue FIRE here). Observe that this solution involves a trivial silent choir, namely, when no agent observes START.

In the presence of byzantine agents, however, this solution does not work, as faulty agents may send a notification without having observed anything. A correct solution for FRR must, hence, prevent the faulty agents from triggering FIRE at any correct agent. In this paper, we will establish the necessary and sufficient state of knowledge for correctly solving FRR in our epistemic reasoning framework for byzantine multi-agent systems [17, 18, 19]. At least since the ground-breaking work by Halpern and Moses [14], the knowledge-based approach [8] has been known as a powerful tool for analyzing distributed systems. In a nutshell, it uses epistemic logic [16] to reason about knowledge and belief in distributed systems. As agents take actions (e.g., FIRE) based on the accumulated local knowledge, reasoning about the latter is useful both for protocol design and impossibility proofs.

In the *runs-and-systems* framework for reasoning about multi-agent systems [8, 14], the set of all possible runs r (executions) of a system I determines the Kripke model, formed by pairs (r, t) of a run $r \in I$ and time $t \in \mathbb{N}$ representing global states $r(t)$. Note that time is modeled as discrete for simplicity, without necessarily being available to the agents. Two pairs (r, t) and (r', t') are indistinguishable for agent i iff i has the same local state in both global states represented by those points, formally, if $r_i(t) = r'_i(t')$. A modal *knowledge operator* K_i is used to capture that agent i knows some fact φ in run $r \in I$ at time $t \in \mathbb{N}$. Formally, $(I, r, t) \models K_i\varphi$ iff for every $r' \in I$ and for every t' with $r_i(t) = r'_i(t')$ it holds that $(I, r', t') \models \varphi$. Note that φ can be a formula containing arbitrary atomic propositions like $\overline{\text{occurred}}(e)$ (event e occurred) or correct_i (i did not fail yet), as well as other knowledge operators and temporal modalities like \diamond (eventually) and \square (always), combined by standard logical operators \neg, \wedge, \vee , and \rightarrow . For example, $(I, r, t) \models \diamond K_i \overline{\text{occurred}}(e)$ states that there is some time $t' \geq t$ when i knows that event e occurred. Important additional modalities for a group G of agents are *mutual knowledge* $E_G\varphi := \bigwedge_{i \in G} K_i\varphi$ and *common knowledge* $C_G\varphi$ that can be informally expressed as an infinite conjunction $C_G\varphi \equiv E_G\varphi \wedge E_G(E_G\varphi) \wedge \dots$; in other words, this means that every agent in G knows φ , and every agent in G knows that every agent in G knows φ , and so on.

Actions performed by the agents when executing a protocol take place when they have accumulated some specific epistemic knowledge. According to the pivotal *Knowledge of Preconditions Principle* [22], it is universally true that if φ is a necessary condition for agent i to take a certain action then i may act only if $K_i\varphi$ is true. For example, in order for agent i to decide on 0 in a binary consensus algorithm, i must know that some agent j has started with initial value $x_j = 0$, i.e., $K_i(\exists j : x_j = 0)$ must hold true. Showing that agents act without having attained the respective necessary knowledge is, hence, a very effective way for proving incorrectness of protocols. Conversely, optimal distributed algorithms can be designed by letting agents act as soon as all respective necessary knowledge has been established. Prominent examples are the protocols based on silent choirs analyzed in [13] and the *unbeatable* consensus protocols introduced in [5], which are not just worst-case optimal but also not strictly dominated w.r.t. termination time by any other protocol in *any* execution.

Related work: The knowledge-based approach has been used for studying several distributed computing problems in systems with uncertainty but no failures. In [3], Ben-Zvi and Moses considered the simple *ordered response* problem in distributed systems, where the agents had to respond to an external START event by executing a special one-shot action FIRE in a given order i_1, i_2, \dots . The authors showed that, in every correct solution, agent i_k has to establish nested knowledge $K_{i_k}K_{i_{k-1}} \dots K_{i_1}\overline{\text{occurred}}(\text{START})$ before it can issue FIRE and that this nested knowledge is also sufficient. In the conference version [1] of [3], the authors also considered the *simultaneous response* problem where all agents had to issue FIRE at the same time. It requires the group G of firing agents to establish common knowl-

edge $C_G \overline{\text{occurred}}(\text{START})$ [14]. This work was later extended to responses that are not simultaneous but tightly coordinated in time [2, 12].

The knowledge-based approach has also been successfully applied to fault-tolerant synchronous distributed systems. Agents suffering from crash or omission failures have been studied in [24, 25], primarily in the context of agreement problems [7, 15], which require some form of common knowledge. Important ingredients here are the indexical set of correct agents and a related belief operator $B_i\varphi := K_i(\text{correct}_i \rightarrow \varphi)$ [23], which states that agent i knows φ to be true in all runs where i is correct. This notion of “defeasible knowledge” also underlies a variant of common knowledge that has been used successfully for characterizing simultaneous distributed agreement [24, 25]. Closer related to our FRR problem is eventual distributed agreement studied in [15], where the stronger notion of continual common knowledge proved its value. The latter needs to hold throughout a run, i.e., from the beginning, which makes sense here since it is only applied to conditions on the initial state. Continual common knowledge does not seem readily applicable to FRR, however, as START can occur at any time in a run. More recent results are the already mentioned unbeatable consensus algorithms in synchronous systems with crash failures [5] and the silent-choir based message-optimal protocols [13].

Detailed contributions: We rigorously define the FRR problem and its weaker variant FR, without the all-or-nothing requirement (agreement), in epistemic terms and identify the necessary and sufficient state of knowledge that must be established by a correct agent in order to issue FIRE in every correct solution for FRR. Since FRR involves distributed agreement, the required state of knowledge involves some form of (eventual) common knowledge of $\overline{\text{occurred}}(\text{START})$. Interestingly, it turned out that establishing the respective eventual mutual knowledge (namely, “eventual mutual hope” where the hope modality is defined as $H_i\varphi := \text{correct}_i \rightarrow B_i\varphi$) already implies the required common knowledge (namely, “eventual common hope”). We also identify subtle variations of the necessary and sufficient state of knowledge for FRR for different assumptions on the occurrence of START.

Whereas identifying the necessary and sufficient state of knowledge for the agents to FIRE does not immediately lead to efficient practical protocols, it is an important first step towards this goal. Indeed, as for the ordered response problem in [3], for example, we expect this knowledge to lead to necessary and sufficient *communication structures*, which must be present in every run of any correct protocol solving FRR. Knowing the latter would not only enable us to decide right away whether the communication guarantees provided by some distributed system allow to solve FRR, but also facilitate the design of efficient protocols.

Paper organization: In Section 2, we introduce some minimal basic notation from our modeling framework [19]. In Section 3, we provide the detailed definition and an epistemic analysis of the FRR problem. Some conclusions and directions of future work in Section 4 complete our paper.

2 Preliminaries

In this section, we outline the basic concepts and facts that are employed in our epistemic analysis of Firing Rebels with Relay (FRR). Our analysis was actually performed within the rigorous framework (that is based on the standard runs-and-systems framework) first developed in [19], which incorporates agents’ ability to arbitrarily deviate from normative behavior. This framework enables one to formally express the epistemic limitations (of the agents) that the presence of possibly fully byzantine agents in the system imposes.

In particular, we proved in [18] that asynchronous agents in a system with fully byzantine agents can never *know* that a particular event took place or even that an agent performed a particular action (since

the local state of a malfunctioning agent may have been corrupted, the above statement applies even to the agent's knowledge of its own actions). This rather disappointing result stems from the inability of even correct agents to exclude the possibility of the so-called *brain-in-a-vat scenario*. In other words, an agent can never be sure that the events and actions recorded in its local history truly happened as recorded rather than being figments of its own malfunction. To make matters worse, agents can also never *know* that they are correct either. Thus, unable to rely on knowledge or their own correctness, our agents are forced to rely instead on *belief* $B_i\varphi := K_i(\text{correct}_i \rightarrow \varphi)$ (cf. [23]). In this paper, we show that even belief is not always appropriate and needs to be replaced with a modality we called *hope* defined as $H_i\varphi := \text{correct}_i \rightarrow B_i\varphi$ (for a detailed explanation see Remark 11).

Since the epistemic analysis presented in this paper is protocol-independent and does not rely on the artefacts of our modeling, we omit all details irrelevant to the task at hand and present our findings in an epistemic language with temporal modalities that is interpreted in Kripke models generated by runs in our framework. The purpose of this section is to provide all the necessary ingredients (referring the reader to [17, 18, 19] for full details of the said framework).

We fix a finite set $\mathcal{A} = \{1, \dots, n\}$ of *asynchronous agents* with *perfect recall*. Each agent $i \in \mathcal{A}$ can perform *actions* (according to its protocol), e.g., send *messages*. One of the actions any agent can perform is FIRE. Agents also witness *events* (triggered by the *environment*) such as message delivery. One of the events that can be observed by any agent is START. We use a discrete time model governed by a global clock with domain $\mathbb{T} = \mathbb{N}$. All events taking place after clock time $t \in \mathbb{T}$ and no later than $t + 1$ are grouped into a *round* denoted $t + \frac{1}{2}$ and are treated as happening simultaneously. Apart from actions, everything in the system is governed by the *environment*. Unlike the environment, agents only have limited local information, in particular, being asynchronous, do not have access to the global clock. This is achieved by allowing them not to perform actions in some rounds and allowing them, in the absence of either actions or events, to stay in the same local state for several rounds in a row. The agents have perfect recall in the sense that, once recorded in their local history, actions and events are never forgotten.

No assumptions apart from liveness are made about the communication. Messages can be lost, arbitrarily delayed, and/or delivered in the wrong order. In addition, the environment may cause at most f agents to become *byzantine* faulty. A byzantine faulty agent can perform any action irrespective of its protocol and “observe” events that did not happen. It can also have false memories about actions it has performed. At the same time, like the global clock, such malfunctions are not directly visible to an agent.

Throughout the paper, horizontal bars signify phenomena that are correct. Note that the absence of this bar should not be equated to faultiness but rather means the absence of a claim of correctness.

Agent i 's local view of the system immediately after round $t + \frac{1}{2}$, referred to as (*process-time* or *agent-time*) *node* $(i, t + 1)$, is recorded in i 's *local state* $r_i(t + 1)$, also called i 's *local history*. A *run* r is a sequence of *global states* $r(t) = (r_\varepsilon(t), r_1(t), \dots, r_n(t))$ of the whole system consisting of the *state* $r_\varepsilon(t)$ of the *environment* and local states $r_i(t)$ of every agent. Unlike local states, the global state of the system necessarily updates every round to include all actions and events that happened (even the empty set thereof is faithfully recorded and modifies the global state). The set of all global states is denoted \mathcal{G} .

What happens in each round is determined by nondeterministic protocols P_i of the agents, the non-deterministic protocol P_ε of the environment, and chance, the latter implemented as the *adversary* part of the environment (the exact technical details are not important for this paper).

In our epistemic analysis, we consider pairs (r, t) of a run r and time t . A *valuation function* π determines whether an atomic proposition from *Prop* is true in run r at time t . The determination is arbitrary except for a small set of *designated atomic propositions* whose truth value at (r, t) is fully determined by the state of the system. More specifically, for $i \in \mathcal{A}$ and $t \in \mathbb{T}$,

correct_i is true at (r, t) iff no faulty event happened to i by time t ;

$\overline{\text{occurred}}_i(o)$ is true at (r,t) iff i 's local history $r_i(t)$ contains an *accurate* record of action/event o occurring (for example, in this paper we use action $o = \text{FIRE}$ and event $o = \text{START}$);

$$\overline{\text{occurred}}(o) := \bigvee_{i \in \mathcal{A}} \overline{\text{occurred}}_i(o).$$

An *interpreted system* is a pair $I = (R, \pi)$ where R is the set of considered runs. The language is $\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_i\varphi \mid \diamond\varphi \mid Y\varphi$ where $p \in \text{Prop}$ and $i \in \mathcal{A}$; derived Boolean connectives are defined in the usual way; $\Box\varphi := \neg\diamond\neg\varphi$. Truth for these *formulas* is defined in the standard way, in particular, for a run $r \in R$, time $t \in \mathbb{T}$, atomic proposition $p \in \text{Prop}$, agent $i \in \mathcal{A}$, and formula φ , we have $(I, r, t) \models p$ iff $(r, t) \in \pi(p)$, and $(I, r, t) \models K_i\varphi$ iff $(I, r', t') \models \varphi$ for any $r' \in R$ and $t' \in \mathbb{T}$ such that $r_i(t) = r'_i(t')$, and $(I, r, t) \models \diamond\varphi$ iff $(I, r, t') \models \varphi$ for some $t' \geq t$, and $(I, r, t) \models Y\varphi$ iff $t > 0$ and $(I, r, t-1) \models \varphi$. A formula φ is valid in I , written $I \models \varphi$, iff $(I, r, t) \models \varphi$ for all $r \in R$ and $t \in \mathbb{T}$.

3 The Firing Rebels Problem

In [27], Srikanth and Toueg introduced the consistent broadcasting primitive, which proved its value in several different contexts, ranging from low-level byzantine fault-tolerant tick generation in various system models [6, 11, 26, 30] to classic clock synchronization [27] to Byzantine Agreement [28]. As argued already in Section 1, it gave rise to our *Firing Rebels with Relay* problem FRR [9], which can be seen as a natural generalization of a *silent choir* in byzantine fault-tolerant systems [13]. As an important building block for designing byzantine fault-tolerant systems, it is therefore a natural target for a detailed epistemic analysis in our framework [19].

Our Firing Rebels problems assume that every agent $i \in \mathcal{A}$ may observe an event START and may generate an action FIRE according to the following specification:

Definition 1 (Firing Rebels with and without Relay). *A system is consistent with Firing Rebels (FR) for $f \geq 0$ when all runs satisfy:*

(C) *Correctness: If at least $2f + 1$ agents learn that START occurred at a correct agent, all correct agents perform FIRE eventually.*

(U) *Unforgeability: If a correct agent performs FIRE, then START occurred at a correct agent.*

Moreover, the system is consistent with Firing Rebels with Relay (FRR) if every run also satisfies:

(R) *Relay: If a correct agent performs FIRE, all correct agents perform FIRE eventually.*

Remark 2 (Variants of Correctness). *A different specification for Correctness can sometimes be found in literature: “If at least $f + 1$ reliable agents locally observed START, then some reliable agent fires eventually” (see, e.g., [4]). Here, a reliable agent is one that will always follow its protocol, which corresponds to a forever correct agent in our terminology. In the case of FRR, by invoking (R), this specification implies “If at least $f + 1$ reliable agents locally observed START, then all reliable agents fire eventually.” Relying on such a specification in asynchronous settings is problematic, however, because reliability depends on the future behavior of the system. Even complete knowledge of the global state, at a given time in a run, does not allow to identify the reliable agents whose observations of START could be relied upon. Thus, we require $2f + 1$ arbitrary (correct or faulty) agents instead. Of course, given the limit of f faulty agents per run, at least $f + 1$ (not necessarily the same) of these agents will remain reliable in every run. Moreover, we relax the condition of the $2f + 1$ agents locally observing START to each of them learning that START happened to some correct agent. This is preferable, because direct observation is only one possible way of ascertaining that START occurred. For instance, if an agent has*

already determined² who the f faulty agents are, e.g., due to their erratic behavior in the past, then a confirmation of START from just one other agent would be sufficient.

We use the following abbreviations:

$$\begin{aligned} B_i\varphi &:= K_i(\text{correct}_i \rightarrow \varphi) & H_i\varphi &:= \text{correct}_i \rightarrow B_i\varphi = \text{correct}_i \rightarrow K_i(\text{correct}_i \rightarrow \varphi) \\ E^B\varphi &:= \bigwedge_{j \in \mathcal{A}} B_j\varphi & E^H\varphi &:= \bigwedge_{j \in \mathcal{A}} H_j\varphi \\ E^{\diamond B}\varphi &:= \bigwedge_{j \in \mathcal{A}} \diamond B_j\varphi & E^{\diamond H}\varphi &:= \bigwedge_{j \in \mathcal{A}} \diamond H_j\varphi \end{aligned}$$

It has been shown in [10] that hope is a normal modality, in particular, $\models H_i(\varphi \wedge \psi) \rightarrow H_i\varphi \wedge H_i\psi$. We define eventual common hope $C^{\diamond H}\varphi$ as the greatest fixed point of the equation $\chi \leftrightarrow E^{\diamond H}(\varphi \wedge \chi)$ in the standard way (using the Knaster–Tarski theorem [29]) and use the following properties (the general versions of which can be found in Lemma 11.5.7 in [8]): for any interpreted system I ,

$$I \models C^{\diamond H}\varphi \leftrightarrow E^{\diamond H}(\varphi \wedge C^{\diamond H}\varphi); \quad (1)$$

$$\text{if } I \models \psi \rightarrow E^{\diamond H}(\varphi \wedge \psi), \quad \text{then } I \models \psi \rightarrow C^{\diamond H}\varphi. \quad (2)$$

3.1 Modeling

Definition 3. For an agent $i \in \mathcal{A}$, we define:

$$\begin{aligned} \overline{\text{start}}_i &:= Y\overline{\text{occurred}}_i(\text{START}) \wedge \text{correct}_i & \overline{\text{fire}}_i &:= \overline{\text{occurred}}_i(\text{FIRE}) \wedge \text{correct}_i \\ \overline{\text{start}} &:= \bigvee_{j \in \mathcal{A}} \overline{\text{start}}_j & \overline{\text{fire}} &:= \bigvee_{j \in \mathcal{A}} \overline{\text{fire}}_j \end{aligned}$$

Note that for one of these formulas to be true, it is necessary for (one of) the involved agent(s) to be correct not only at the time the event/action in question occurred but also at the time of the evaluation. Using the yesterday modality Y in $\overline{\text{start}}_i$ accounts for the fact that agents cannot act on a precondition in the same round it is established.

Using Def. 3, we can translate the specification of FRR (stated in Def. 1) as follows:

Definition 4 (Modeling Firing Rebels). *An interpreted system I is consistent with Firing Rebels with Relay for $f \geq 0$ if the following conditions Correctness (C), Unforgeability (U), and Relay (R) hold:*

$$\begin{aligned} \text{(C)} \quad I &\models \bigvee_{\substack{G \subseteq \mathcal{A} \\ |G|=2f+1}} \bigwedge_{j \in G} K_j(\text{correct}_j \rightarrow \overline{\text{start}}) \rightarrow \bigwedge_{i \in \mathcal{A}} \diamond(\text{correct}_i \rightarrow \overline{\text{fire}}_i) \\ \text{(U)} \quad I &\models \overline{\text{fire}} \rightarrow \overline{\text{start}} \\ \text{(R)} \quad I &\models \overline{\text{fire}} \rightarrow \bigwedge_{i \in \mathcal{A}} \diamond(\text{correct}_i \rightarrow \overline{\text{fire}}_i) \end{aligned}$$

Remark 5 (Variants of eventuality). *The phrase all correct agents fulfill φ_i eventually in Def. 1 can be formalized in two different ways:*

- $\bigwedge_{i \in \mathcal{A}} \diamond(\text{correct}_i \rightarrow \varphi_i)$ states that each agent will either become faulty at some point in the future or will fulfill its respective φ_i at some point in the future.

²Strictly speaking, the agent in this situation *does not know* that the f agents are faulty, but rather that they are faulty if it itself is not. By the same token, whenever we say “learned,” “determined,” or “ascertained” above, what we mean is reasoning under the assumption of its own correctness, i.e., the belief modality B_i rather than the knowledge modality K_i .

- $\diamond \bigwedge_{i \in \mathcal{A}} (\text{correct}_i \rightarrow \varphi_i)$ states that there is one moment in the future by which every agent still correct fulfills its respective φ_i .

The second statement is a strengthening of the first by demanding agents to have one common moment by which all correct agents fulfill their respective φ_i 's. On the other hand, the first variant is a *more intuitive* reading that is more widely applicable. Fortunately, for our model of FRR with $\varphi_i = \overline{\text{fire}}_i$, the two formulations are equivalent because, due to agents having perfect recall, $\text{correct}_i \rightarrow \overline{\text{fire}}_i$ is a stable fact.

3.2 Necessary and Sufficient Level of Knowledge

The goal of this subsection is to

- strengthen the given necessary conditions on a single agent's firing — namely, that $\overline{\text{start}}$ must hold by Unforgeability (U) and $\bigwedge_{i \in \mathcal{A}} \diamond (\text{correct}_i \rightarrow \overline{\text{fire}}_i)$ must hold by Relay (R) — to statements that describe the state of knowledge necessary for the agent to achieve before firing;
- show that firing upon reaching this state of knowledge is sufficient for satisfying the conditions Unforgeability (U) and Relay (R) on all correct agents eventually firing;
- show how Correctness (C) helps simplify these necessary and sufficient conditions in the presence of sufficiently many agents.

Thus, protocols prescribing an agent to fire as soon as this state of knowledge is achieved are correct and optimal in the sense that firing earlier would violate the necessary conditions whereas firing prescribed by this state of knowledge is guaranteed to fulfill all requirements of FRR.

Note that the case when insufficiently many agents learn that START occurred at a correct agent trivially satisfies condition (C). In this case, FRR reduces to (U)+(R), a problem with a trivial solution of all correct agents not firing. It is the combination of all all three conditions that makes FRR a problem worth the analysis.

The first lemma formalizes the fact that, since agents have perfect recall of their past perceptions, reasoning under the assumption of their own correctness leads them to *believe* that these perceptions were accurate. For instance, an agent who recalls observing START believes that, unless it is faulty, a correct agent (namely, itself) observed START. (A formal proof can be found in the Appendix on p. 151.)

Lemma 6. *For any interpreted system I and any agent $i \in \mathcal{A}$:*

$$I \models \overline{\text{fire}}_i \rightarrow B_i \overline{\text{fire}}_i \quad (3)$$

$$I \models \overline{\text{fire}}_i \rightarrow B_i \overline{\text{fire}} \quad (4)$$

$$I \models \overline{\text{start}}_i \rightarrow B_i \overline{\text{start}}_i \quad (5)$$

$$I \models \overline{\text{start}}_i \rightarrow B_i \overline{\text{start}} \quad (6)$$

Unforgeability (U) states that $\overline{\text{start}}$ is a necessary condition for a correct agent firing. It follows from the Knowledge of Preconditions Principle that any correct agent must ascertain $\overline{\text{start}}$ (modulo its own correctness) before firing. We formalize this argument and provide an independent proof:

Lemma 7 (State of knowledge necessary for firing in presence of Unforgeability (U)). *Let I be an interpreted system consistent with Unforgeability (U). For any agent $i \in \mathcal{A}$,*

$$I \models \overline{\text{fire}}_i \rightarrow B_i \overline{\text{start}}. \quad (7)$$

Proof. Immediately follows from (4), (U), and monotonicity/normality of B_i . □

Corollary 8. *For any interpreted system consistent with FR, (7) is satisfied for all agents.*

Similarly, lifting the Relay condition (R) to the level of agent's knowledge yields the requirement that, in order to fire, an agent must believe that all correct agents eventually will have fired.

Lemma 9 (State of knowledge necessary for firing in presence of Relay (R)). *Let I be an interpreted system consistent with Relay (R). For any agent $i \in \mathcal{A}$,*

$$I \models \overline{fire}_i \rightarrow B_i \bigwedge_{j \in \mathcal{A}} \diamond(\overline{correct}_j \rightarrow \overline{fire}_j). \quad (8)$$

Proof. Immediately follows from (4), (R), and monotonicity of B_i . \square

Combining the conditions necessary for (U) and (R), we establish the following level of knowledge necessary for firing in FRR (a proof can be found in the Appendix on p. 151):

Theorem 10 (State of knowledge necessary for firing in presence of both (U) and (R)). *Let I be an interpreted system consistent with (U) and (R). For any agent $i \in \mathcal{A}$,*

$$I \models \overline{fire}_i \rightarrow B_i \left(\overline{start} \wedge E^{\diamond H} \overline{start} \right).$$

Remark 11 (Emergence of hope). *Note that $I \models \overline{fire}_i \rightarrow B_i E^{\diamond B} \overline{start}$ does not generally hold. We cannot strengthen the necessary condition by replacing eventual mutual hope with eventual mutual belief, i.e., by omitting $\overline{correct}_j$ therein. In other words, the use of hope for deeper iterations of knowledge modalities is crucial for the correct formulation. Indeed, in the case of our notion of belief, agent i can rarely have unconditional beliefs about another agent j 's beliefs. The problematic situation is when agent j 's perception is compromised. In that case, agent i has no way of ascertaining what j 's erroneous input data might be and, hence, cannot determine what a correct agent would have inferred from these incorrect inputs. According to our notion of belief, whether agent i itself is correct or not, it reasons assuming that its own perceptions are the objective reality. The $\overline{correct}_j$ assumption is, therefore, necessary to anchor j to the same (allegedly) objective reality contemplated by i , even though j 's access to the facts of this objective reality is generally different from i 's. Note also that j 's reasoning is generally happening in the future relative to i 's current reasoning, meaning that we also implicitly assume reality to be stable.*

Remark 12 (Relation to indexical sets). *Another approach to describing beliefs of fault-prone agents is via so-called indexical sets [8, 25], which are variable (non-rigid) sets that can be used to represent the set of all correct agents at every point in the system. While our results could be reformulated in terms of indexical sets, there were several reasons for us to choose another language. Besides the ability to reason about all agents, whether correct or faulty, in a uniform way, we tried to stay as close as possible to the standard language of epistemic modal logic. Perhaps more importantly, however, was the moral lesson of the already mentioned Knowledge of Preconditions Principle [22], which reveals how important it is for an agent to know all ingredients affecting its behavior, correctness of itself and other agents being one of them. Thus, we believe that the transparent and explicit use of correctness in our language is advantageous. An immediate example is the distinction between belief and hope discussed in Remark 11, which would have remained somewhat obscured in the indexical set notation.*

Remark 13 (Eventual mutual hope is not sufficient). *While using $B_i (\overline{start} \wedge E^{\diamond H} \overline{start})$ as a trigger for agent i firing will ensure Unforgeability (U), it is too weak to guarantee Relay (R). Indeed, consider a system with 3 agents ($n = 3$), at most one of which can become faulty ($f = 1$). In such a system, receiving the same information from two independent sources is sufficient to believe in its validity, while information from only one source without observing it first hand is not. Suppose that the protocol forces*

a correct agent to notify all other agents whenever it observed START. Consider a run where agent b is byzantine from the beginning, whereas agents c_1 and c_2 remain correct. Let c_1 and c_2 each observe START and, hence, notify all agents about it. Meanwhile b falsely notifies c_2 that it too observed START but will never duplicate this message to c_1 . Thus,

- correct c_2 observed START and eventually received 2 confirmations of START from c_1 and b ;
- correct c_1 observed START and eventually received 1 confirmation of START from c_2 ;
- faulty b did not observe START but was eventually notified of START by both c_1 and c_2 .

In this situation, all agents eventually believe that START was correctly observed (c_1 and c_2 saw it themselves, whereas b has 2 independent confirmations). Moreover, c_2 has a reason to believe in the eventual mutual hope of START. Indeed, hope would be trivially satisfied for a faulty agent, whereas any correct agent would eventually receive at least 2 confirmations out of 3 that c_2 itself possesses. Thus, according to the proposed knowledge threshold, c_2 should fire. On the other hand, c_1 will never fire because it cannot be sure that b will eventually hope that START occurred. In c_1 's mind, if b were correct and c_2 were faulty and did not send a confirmation to b , then b would only ever receive 1 confirmation, which is not sufficient to make it trust START truly occurred. Hence, c_1 would never fire, and Relay (R) would be violated.

The issue here is that $B_i E^{\diamond H} \overline{\text{start}}$ for one correct agent i does not generally imply that eventually $B_j E^{\diamond H} \overline{\text{start}}$ for all other correct agents j .

Thus, although $B_i E^{\diamond H} \overline{\text{start}}$ is necessary before i can fire and is in principle actionable, acting on it may be premature. The necessary state of knowledge must be further strengthened. Since FRR involves an agreement property (one correct agent fires only if all other correct agents also fire eventually), it is not very surprising that, in fact, some form of common knowledge, specifically *eventual common hope*, plays a role. We have shown (see a proof in the Appendix on p. 152) that Unforgeability and Relay together imply that, in order to fire an agent must ascertain (modulo its own correctness) both that START was observed by some correct agent and the eventual common hope of the same fact:

Theorem 14 (State of knowledge necessary for firing in presence of both (U) and (R)). *Let I be an interpreted system consistent with (U) and (R). For any agent $i \in \mathcal{A}$,*

$$I \models \overline{\text{fire}}_i \rightarrow B_i \left(\overline{\text{start}} \wedge C^{\diamond H} \overline{\text{start}} \right). \quad (9)$$

Corollary 15. *For any interpreted system consistent with FRR, (9) is satisfied for all agents.*

We now show (see a proof in the Appendix on p. 152) that, unlike belief in eventual mutual hope (see Remark 13), belief in eventual *common hope* is sufficient to fulfill Unforgeability and Relay, i.e., that firing as soon as the necessary state of knowledge from Theorem 14 is achieved does guarantee that both (U) and (R) are fulfilled:

Theorem 16 (Sufficient conditions for (U) and (R)). *For any interpreted system I :*

1. (U) is fulfilled if $I \models \bigwedge_{i \in \mathcal{A}} (\neg B_i \overline{\text{start}} \rightarrow \neg \overline{\text{fire}}_i)$.
2. Both (U) and (R) are fulfilled if

$$I \models \bigwedge_{i \in \mathcal{A}} \left(\left(\neg B_i \left(\overline{\text{start}} \wedge C^{\diamond H} \overline{\text{start}} \right) \rightarrow \neg \overline{\text{fire}}_i \right) \wedge \left(B_i \left(\overline{\text{start}} \wedge C^{\diamond H} \overline{\text{start}} \right) \rightarrow \diamond (\text{correct}_i \rightarrow \overline{\text{fire}}_i) \right) \right). \quad (10)$$

Remark 17 (Belief in $\overline{\text{start}}$ is not redundant). *Since common knowledge is the strongest type of knowledge and knowledge is supposed to be factive, it might be tempting to think that the conjunct $\overline{\text{start}}$ is redundant in the formulations of Theorems 10, 14, and 16.2. The difference in our setting is that the relevant epistemic state is eventual, meaning that it need not be factual at present. Still, one might question how it could be possible to achieve even an eventual knowledge/belief/hope without the event actually happening. Indeed, if there is no reason for agents to expect START to necessarily occur, their predictions about START occurring can only rely on it already having occurred. This observation is formalized in Lemma 23 and the immediately following corollary.*

Definition 18 (Potentially persistent formulas). *A formula φ is called potentially persistent in an interpreted system $I = (R, \pi)$ if, for any run $r \in R$ and any time $t \in \mathbb{T}$ such that $(I, r, t) \models \varphi$, there exists a run $r' \in R$ such that $r'(t) = r(t)$ — i.e., r' is an alternative continuation of the global state $r(t)$ — and such that $(I, r', t) \models \Box\varphi$. In other words, a true potentially persistent formula can stay true forever.*

The following 3 lemmas follow from definitions (proofs of the last two are in the Appendix on p. 152).

Lemma 19. $I \models \neg \text{correct}_i \rightarrow \Box \neg \text{correct}_i$ for any $i \in \mathcal{A}$ and interpreted system I .

Lemma 20. $I \models K_i \Diamond \neg \varphi \rightarrow K_i \neg \varphi$ for any $i \in \mathcal{A}$ and φ potentially persistent in an interpreted system I .

Lemma 21. $I \models B_i \Diamond (\text{correct}_i \rightarrow \varphi) \leftrightarrow K_i \Diamond (\text{correct}_i \rightarrow \varphi)$ for any $i \in \mathcal{A}$, formula φ , and interpreted system I , i.e., believing something eventually happens modulo one's own correctness is as strong as knowing it eventually happens modulo one's own correctness.

Corollary 22. $I \models B_i \Diamond H_i \varphi \leftrightarrow K_i \Diamond H_i \varphi$ for any $i \in \mathcal{A}$, formula φ , and interpreted system I .

Lemma 23 (Early local belief). *If $\text{correct}_i \wedge \overline{\text{start}}$ is potentially persistent in an interpreted system I ,*

$$I \models B_i \Diamond H_i \overline{\text{start}} \rightarrow B_i \overline{\text{start}}.$$

Proof. A proof can be found in the Appendix on p. 153. □

Noting that $I \models C^{\Diamond H} \overline{\text{start}} \rightarrow E^{\Diamond H} \overline{\text{start}}$ because of the normality of the hope modality [10], we can derive:

Corollary 24. *If $\text{correct}_i \wedge \overline{\text{start}}$ is potentially persistent in an interpreted system I , then*

$$I \models B_i E^{\Diamond H} \overline{\text{start}} \rightarrow B_i \overline{\text{start}}, \quad (11)$$

$$I \models B_i C^{\Diamond H} \overline{\text{start}} \rightarrow B_i \overline{\text{start}}. \quad (12)$$

Remark 25 (Conditions on dropping $B_i \overline{\text{start}}$). *If, contrary to the conditions of the early local belief lemma, $\overline{\text{start}}$ is inevitable, agents may be able to predict the eventual arrival of START before the fact. For instance, if START eventually happens to every agent, i.e., if $I \models \Diamond \text{occurred}_i(\text{START})$, then $I \models \Diamond B_i \overline{\text{start}}$ for every agent $i \in \mathcal{A}$. It follows that $I \models CC^{\Diamond H} \overline{\text{start}}$, i.e., it is common knowledge, from the very beginning, that there is eventual common hope of $\overline{\text{start}}$. Thus, this state of knowledge is achieved independently of START happening, and triggering FIRE risks violating Unforgeability (U).*

On the other hand, even if START is assured, it may not always be possible to predict it in advance. While sufficient for dropping the conjunct $\overline{\text{start}}$ from the conditions triggering FIRE in Theorem 16.2, the potential persistency of $\text{correct}_i \wedge \overline{\text{start}}$ is not necessary. Indeed, (12) can hold even when START is always guaranteed to happen. For instance, in an interpreted system where START happens exactly once per run, no agent ever becomes faulty, and, in addition, agents never communicate, $I \models \neg B_i E^{\Diamond H} \overline{\text{start}}$

because only the agent who observed START can learn that it already occurred. All the others can only be sure that START will occur eventually. By (1), also $I \models \neg B_i C^{\diamond H} \overline{start}$. Thus, both implications (11) and (12) are vacuously true, allowing to drop \overline{start} without affecting the behavior of agents, though admittedly in such interpreted systems agents should never fire anyways.

The following ‘‘Lifting Lemma’’ shows that Correctness (C) lifts eventual mutual hope to eventual common hope. This way, the arbitrarily deep nested hope implied by the latter effectively collapses, a phenomenon that has also been reported for other problems [1]. A proof of the lemma can be found in the Appendix on p. 153.

Lemma 26 (Lifting Lemma). *Let I be an interpreted system consistent with (C) and let $|\mathcal{A}| \geq 3f + 1$, where $f \geq 0$ is the maximum number of byzantine faulty agents in a run. Furthermore, assume that*

$$I \models \overline{fire}_i \rightarrow B_i \left(\overline{start} \wedge E^{\diamond H} \overline{start} \right) \quad (13)$$

holds. Then,

$$I \models E^{\diamond H} \overline{start} \rightarrow C^{\diamond H} \overline{start}. \quad (14)$$

Corollary 27. *Let I be an interpreted system with at least $3f + 1$ agents. If I is consistent with FRR, then (14) holds.*

Proof. In interpreted systems consistent with (U) and (R), property (13) follows from Theorem 10. \square

4 Conclusions and Future Work

We introduced a problem called Firing Rebels with Relay (FRR) and its weaker variant called Firing Rebels (FR), which capture the essentials of a well-known building block for byzantine fault-tolerant distributed algorithms. The main purpose of our paper was to determine the state of knowledge correct agents must achieve in order to act (FIRE) according to the specification of the problem at hand. Through a detailed epistemic analysis, we established that the necessary and sufficient levels of knowledge required for acting rely on the novel notion of eventual common hope. We also found the conditions under which a single level of eventual mutual hope can guarantee infinitely many levels of eventual common hope and explored the surprisingly non-trivial relationship of the eventual common hope of START with the actual appearance of START.

Regarding future work, our next step is to complete the characterization of (eventual) common hope. More precisely, what remains to be done is developing an independent axiomatization of the (eventual) common hope modality based on our existing axiomatization of the hope modality (which does not depend on the knowledge modality). In addition, we are working on identifying necessary and sufficient communication structures and optimal protocols for FRR.

References

- [1] Ido Ben-Zvi & Yoram Moses (2010): *Beyond Lamport’s Happened-Before: On the Role of Time Bounds in Synchronous Systems*. In Nancy A. Lynch & Alexander A. Shvartsman, editors: *DISC 2010: Distributed Computing*, LNCS 6343, Springer, pp. 421–436, doi:10.1007/978-3-642-15763-9_42.
- [2] Ido Ben-Zvi & Yoram Moses (2013): *Agent-Time Epistemics and Coordination*. In Kamal Lodaya, editor: *ICLA 2013: Logic and Its Applications*, LNCS 7750, Springer, pp. 97–108, doi:10.1007/978-3-642-36039-8_9.

- [3] Ido Ben-Zvi & Yoram Moses (2014): *Beyond Lamport's Happened-before: On Time Bounds and the Ordering of Events in Distributed Systems*. *Journal of the ACM* 61(2:13), doi:10.1145/2542181.
- [4] James E. Burns & Nancy A. Lynch (1987): *The Byzantine Firing Squad Problem*. In Franco P. Preparata, editor: *Parallel and Distributed Computing, Advances in Computing Research: A research annual 4*, JAI Press, pp. 147–161. Available at <https://apps.dtic.mil/docs/citations/ADA154770>.
- [5] Armando Castañeda, Yannai A. Gonczarowski & Yoram Moses (2014): *Unbeatable Consensus*. In Fabian Kuhn, editor: *DISC 2014: Distributed Computing, LNCS 8784*, Springer, pp. 91–106, doi:10.1007/978-3-662-45174-8_7.
- [6] Danny Dolev, Matthias Függer, Markus Posch, Ulrich Schmid, Andreas Steininger & Christoph Lenzen (2014): *Rigorously modeling self-stabilizing fault-tolerant circuits: An ultra-robust clocking scheme for systems-on-chip*. *Journal of Computer and System Sciences* 80(2), pp. 860–900, doi:10.1016/j.jcss.2014.01.001.
- [7] Cynthia Dwork & Yoram Moses (1990): *Knowledge and Common Knowledge in a Byzantine Environment: Crash Failures*. *Information and Computation* 88(2), pp. 156–186, doi:10.1016/0890-5401(90)90014-9.
- [8] Ronald Fagin, Joseph Y. Halpern, Yoram Moses & Moshe Y. Vardi (1995): *Reasoning About Knowledge*. MIT Press.
- [9] Patrik Fimml (2018): *Temporal-Epistemic Logic in Byzantine Message-Passing Contexts*. Master's thesis, Technische Universität Wien, Institut für Computer Engineering. Available at http://publik.tuwien.ac.at/files/publik_273448.pdf.
- [10] Krisztina Fruzsa (2019): *Hope for Epistemic Reasoning with Faulty Agents!* In: *ESSLLI 2019 Student Session, FOLLI*. Available at http://esslli2019.folli.info/wp-content/uploads/2019/08/tentative_proceedings.pdf.
- [11] Matthias Függer & Ulrich Schmid (2012): *Reconciling fault-tolerant distributed computing and systems-on-chip*. *Distributed Computing* 24(6), pp. 323–355, doi:10.1007/s00446-011-0151-7.
- [12] Yannai A. Gonczarowski & Yoram Moses (2013): *Timely Common Knowledge: Characterising Asymmetric Distributed Coordination via Vectorial Fixed Points*. In Burkhard C. Schipper, editor: *TARK 2013: Theoretical Aspects of Rationality and Knowledge*, pp. 79–93. Available at http://www.tark.org/proceedings/tark_jan7_13/p79-gonczarowski.pdf.
- [13] Guy Goren & Yoram Moses (2018): *Silence*. In: *PODC 2018: Principles of Distributed Computing, ACM*, pp. 285–294, doi:10.1145/3212734.3212768.
- [14] Joseph Y. Halpern & Yoram Moses (1990): *Knowledge and Common Knowledge in a Distributed Environment*. *Journal of the ACM* 37(3), pp. 549–587, doi:10.1145/79147.79161.
- [15] Joseph Y. Halpern, Yoram Moses & Orli Waarts (2001): *A characterization of eventual Byzantine agreement*. *SIAM Journal on Computing* 31(3), pp. 838–865, doi:10.1137/S0097539798340217.
- [16] Jaakko Hintikka (1962): *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Cornell University Press.
- [17] Roman Kuznets, Laurent Proserpi, Ulrich Schmid & Krisztina Fruzsa (2019): *Causality and Epistemic Reasoning in Byzantine Multi-Agent Systems*. In Lawrence S. Moss, editor: *TARK 2019: Theoretical Aspects of Rationality and Knowledge, EPTCS 297*, Open Publishing Association, pp. 293–312, doi:10.4204/EPTCS.297.19.
- [18] Roman Kuznets, Laurent Proserpi, Ulrich Schmid & Krisztina Fruzsa (2019): *Epistemic Reasoning with Byzantine-Faulty Agents*. In Andreas Herzig & Andrei Popescu, editors: *FroCoS 2019: Frontiers of Combining Systems, LNCS 11715*, Springer, pp. 259–276, doi:10.1007/978-3-030-29007-8_15.
- [19] Roman Kuznets, Laurent Proserpi, Ulrich Schmid, Krisztina Fruzsa & Lucas Gréaux (2019): *Knowledge in Byzantine Message-Passing Systems I: Framework and the Causal Cone*. Technical Report TUW-260549, TU Wien. Available at https://publik.tuwien.ac.at/files/publik_260549.pdf.

- [20] Leslie Lamport (1978): *Time, Clocks, and the Ordering of Events in a Distributed System*. *Communications of the ACM* 21(7), pp. 558–565, doi:10.1145/359545.359563.
- [21] Leslie Lamport, Robert Shostak & Marshall Pease (1982): *The Byzantine Generals Problem*. *ACM Transactions on Programming Languages and Systems* 4(3), pp. 382–401, doi:10.1145/357172.357176.
- [22] Yoram Moses (2015): *Relating Knowledge and Coordinated Action: The Knowledge of Preconditions Principle*. In R. Ramanujam, editor: *TARK 2015: Theoretical Aspects of Rationality and Knowledge*, EPTCS 215, Open Publishing Association, pp. 231–245, doi:10.4204/EPTCS.215.17.
- [23] Yoram Moses & Yoav Shoham (1993): *Belief as defeasible knowledge*. *Artificial Intelligence* 64(2), pp. 299–321, doi:10.1016/0004-3702(93)90107-M.
- [24] Yoram Moses & Mark R. Tuttle (1986): *Programming Simultaneous Actions Using Common Knowledge: Preliminary Version*. In: *27th Annual Symposium on Foundations of Computer Science*, IEEE, pp. 208–221, doi:10.1109/SFCS.1986.46.
- [25] Yoram Moses & Mark R. Tuttle (1988): *Programming Simultaneous Actions Using Common Knowledge*. *Algorithmica* 3, pp. 121–169, doi:10.1007/BF01762112.
- [26] Peter Robinson & Ulrich Schmid (2011): *The Asynchronous Bounded-Cycle model*. *Theoretical Computer Science* 412(40), pp. 5580–5601, doi:10.1016/j.tcs.2010.08.001.
- [27] T. K. Srikanth & Sam Toueg (1987): *Optimal Clock Synchronization*. *Journal of the ACM* 34(3), pp. 626–645, doi:10.1145/28869.28876.
- [28] T. K. Srikanth & Sam Toueg (1987): *Simulating authenticated broadcasts to derive simple fault-tolerant algorithms*. *Distributed Computing* 2(2), pp. 80–94, doi:10.1007/BF01667080.
- [29] Alfred Tarski (1955): *A lattice-theoretical fixpoint theorem and its applications*. *Pacific Journal of Mathematics* 5(2), pp. 285–309, doi:10.2140/pjm.1955.5.285.
- [30] Josef Widder & Ulrich Schmid (2009): *The Theta-Model: achieving synchrony without clocks*. *Distributed Computing* 22(1), pp. 29–47, doi:10.1007/s00446-009-0080-x.

Appendix

Proof of Lemma 6. The argument is the same for FIRE and START. We only provide it for the former. Let $I = (R, \pi)$. Consider a run $r \in R$ and a node $(i, t) \in \mathcal{A} \times \mathbb{T}$. Assume $(I, r, t) \models \overline{\text{fire}}_i$. Since i has perfect recall and this was a correct FIRE action, it was recorded and still remains in i 's local history $r(t)$. Consider any $r' \in R$ and $t' \in \mathbb{N}$ such that $r_i(t) = r'_i(t')$. Then $r'_i(t')$ also contains a record of FIRE. If $(I, r', t') \models \text{correct}_i$, this record must correspond to a correct action and, consequently, $(I, r', t') \models \overline{\text{fire}}_i$. Since $(I, r', t') \models \text{correct}_i \rightarrow \overline{\text{fire}}_i$ whenever $r_i(t) = r'_i(t')$, we have $(I, r, t) \models K_i(\text{correct}_i \rightarrow \overline{\text{fire}}_i)$, i.e., $(I, r, t) \models B_i \overline{\text{fire}}_i$. The other statement about FIRE follows from $\models \overline{\text{fire}}_i \rightarrow \overline{\text{fire}}$ and the monotonicity/normality of B_i . \square

Proof of Theorem 10. Since the system is consistent with (U), (7) holds according to Lemma 7. Thus, given that B_i is a normal modality, it only remains to show that

$$I \models \overline{\text{fire}}_i \rightarrow B_i E^{\diamond H} \overline{\text{start}}. \quad (15)$$

Since the system is consistent with (R), (8) holds according to Lemma 9. Using the replacement property for positive subformulas and the already discussed validity $I \models \overline{\text{fire}}_j \rightarrow B_j \overline{\text{start}}$ from (7), we obtain $I \models \overline{\text{fire}}_i \rightarrow B_i \bigwedge_{j \in \mathcal{A}} \diamond (\text{correct}_j \rightarrow B_j \overline{\text{start}})$, in other words, (15). \square

Proof of Theorem 14. Since (7) holds by Lemma 7, it is sufficient to demonstrate $I \models \overline{\text{fire}}_i \rightarrow B_i C^{\diamond H} \overline{\text{start}}$. Combining (R) with (4) by applying the replacement property for positive subformulas, we obtain $I \models \overline{\text{fire}} \rightarrow E^{\diamond H} \overline{\text{fire}}$. Thus, using (2) with $\varphi = \psi = \overline{\text{fire}}$, we conclude $I \models \overline{\text{fire}} \rightarrow C^{\diamond H} \overline{\text{fire}}$. Since the greatest fixed point of a monotone operator is itself monotone, it follows from (U) that $I \models \overline{\text{fire}} \rightarrow C^{\diamond H} \overline{\text{start}}$. It remains to use (4) and monotonicity of B_i to obtain $I \models \overline{\text{fire}}_i \rightarrow B_i C^{\diamond H} \overline{\text{start}}$. \square

Proof of Theorem 16. For either assumption, $I \models \overline{\text{fire}}_i \rightarrow B_i \overline{\text{start}}$. Since

$$I \models \overline{\text{fire}}_i \rightarrow \text{correct}_i \quad \text{and} \quad I \models \text{correct}_i \rightarrow (B_i \varphi \rightarrow \varphi) \quad \text{for any formula } \varphi, \quad (16)$$

we have $I \models \overline{\text{fire}}_i \rightarrow \overline{\text{start}}$ for each $i \in \mathcal{A}$. Since $\overline{\text{fire}}$ is $\bigvee_{i \in \mathcal{A}} \overline{\text{fire}}_i$, (U) holds by propositional reasoning.

It remains to show that Relay (R) holds under the assumption of (10). Once again, it is sufficient to demonstrate that, for each $i \in \mathcal{A}$,

$$I \models \overline{\text{fire}}_i \rightarrow \bigwedge_{j \in \mathcal{A}} \diamond(\text{correct}_j \rightarrow \overline{\text{fire}}_j). \quad (17)$$

It follows from the first conjunct of (10) that $I \models \overline{\text{fire}}_i \rightarrow B_i C^{\diamond H} \overline{\text{start}}$. Using (16) again, we conclude that $I \models \overline{\text{fire}}_i \rightarrow C^{\diamond H} \overline{\text{start}}$. Since $I \models C^{\diamond H} \varphi \rightarrow \bigwedge_{j \in \mathcal{A}} \diamond H_j(\varphi \wedge C^{\diamond H} \varphi)$ for any formula φ according to (1),

$$I \models \overline{\text{fire}}_i \rightarrow \bigwedge_{j \in \mathcal{A}} \diamond(\text{correct}_j \rightarrow B_j(\overline{\text{start}} \wedge C^{\diamond H} \overline{\text{start}})). \quad (18)$$

Using the second conjunct of (10) and monotonicity of B_j and \diamond in (18), we obtain

$$I \models \overline{\text{fire}}_i \rightarrow \bigwedge_{j \in \mathcal{A}} \diamond(\text{correct}_j \rightarrow \diamond(\text{correct}_j \rightarrow \overline{\text{fire}}_j)).$$

To get (17), it remains to note that $I \models \diamond(\varphi \rightarrow \diamond(\varphi \rightarrow \psi)) \rightarrow \diamond(\varphi \rightarrow \psi)$ for all formulas φ and ψ . \square

Proof of Lemma 20. Let $I = (R, \pi)$. Assume that $(I, r, t) \not\models K_i \neg \varphi$ for some $r \in R$ and $t \in \mathbb{T}$. Then there exists another run $r' \in R$ and time $t' \in \mathbb{T}$ such that $r_i(t) = r'_i(t')$ and $(I, r', t') \models \varphi$. By the potential persistence of φ , there exists an alternative continuation $r'' \in R$ of the prefix $r'(t')$ such that $r''(t') = r'(t')$ and $(I, r'', t') \models \square \varphi$. Thus, $(I, r'', t') \not\models \diamond \neg \varphi$. It remains to note that $r''_i(t') = r'_i(t') = r_i(t)$. Hence, $(I, r, t) \not\models K_i \diamond \neg \varphi$. \square

Proof of Lemma 21. The right-to-left direction is trivial. Hence, we prove the implication from left to right. Firstly, $\neg \text{correct}_i \rightarrow (\text{correct}_i \rightarrow \varphi)$ is a propositional tautology. Hence,

$$I \models \square \neg \text{correct}_i \rightarrow \square(\text{correct}_i \rightarrow \varphi).$$

Using Lemma 19, the fact that $I \models \square \psi \rightarrow \diamond \psi$ by seriality of temporal modalities, and knowledge necessitation, we obtain

$$I \models K_i(\neg \text{correct}_i \rightarrow \diamond(\text{correct}_i \rightarrow \varphi)).$$

By epistemically internalized propositional reasoning,

$$I \models K_i(\text{correct}_i \rightarrow \diamond(\text{correct}_i \rightarrow \varphi)) \wedge K_i(\neg \text{correct}_i \rightarrow \diamond(\text{correct}_i \rightarrow \varphi)) \rightarrow K_i \diamond(\text{correct}_i \rightarrow \varphi).$$

Since we have just shown the second conjunct above to be valid, we obtain the desired

$$I \models K_i(\text{correct}_i \rightarrow \diamond(\text{correct}_i \rightarrow \varphi)) \rightarrow K_i \diamond(\text{correct}_i \rightarrow \varphi). \quad \square$$

Proof of Lemma 23. By Corollary 22, $I \models B_i \diamond H_i \overline{start} \rightarrow K_i \diamond H_i \overline{start}$. Applying factivity of knowledge and propositional reasoning to the expanded version of $K_i \diamond H_i \overline{start}$ yields

$$I \models K_i \diamond (correct_i \rightarrow K_i (correct_i \rightarrow \overline{start})) \rightarrow K_i \diamond (correct_i \rightarrow \overline{start}).$$

Since $correct_i \wedge \neg \overline{start}$ is potentially persistent, and its negation is equivalent to $correct_i \rightarrow \overline{start}$, we have by Lemma 20 that

$$I \models K_i \diamond (correct_i \rightarrow \overline{start}) \rightarrow K_i (correct_i \rightarrow \overline{start})$$

Combining all implications together, we conclude that $I \models B_i \diamond H_i \overline{start} \rightarrow B_i \overline{start}$. \square

Proof of Lemma 26. Let $I = (R, \pi)$. Assume $(I, r, t) \models E^{\diamond H} \overline{start}$ for some $r \in R$ and time $t \in \mathbb{T}$. This means that, for every agent $j \in \mathcal{A}$, there is some $t'_j \geq t$ such that $(I, r, t'_j) \models H_j \overline{start}$. Since $|\mathcal{A}| \geq 3f + 1$, it follows that there exists a group G of $2f + 1$ correct agents such that $(I, r, t'_j) \models H_j \overline{start}$ for all $j \in G$. Since these agents are correct,³ we have $(I, r, t'_j) \models B_j \overline{start}$, i.e., $(I, r, t'_j) \models K_j (correct_j \rightarrow \overline{start})$ for all $j \in G$. Let $t' := \max\{t'_j \mid j \in G\}$. We claim that

$$(I, r, t') \models \bigwedge_{j \in G} K_j (correct_j \rightarrow \overline{start}). \quad (19)$$

Indeed, for any agent $j \in G$ consider any alternative run $\bar{r} \in R$ and time $\bar{t}' \in \mathbb{T}$ such that $\bar{r}_j(\bar{t}') = r_j(t')$. Given that $t' \geq t'_j$ and our agents have perfect recall, there must exist some time $\bar{t}'_j \leq \bar{t}'$ such that $\bar{r}_j(\bar{t}'_j) = r_j(t'_j)$. Thus, $(I, \bar{r}, \bar{t}'_j) \models correct_j \rightarrow \overline{start}$. Since the latter formula is stable, it remains true in \bar{r} by the time \bar{t}' . We showed that $(I, \bar{r}, \bar{t}') \models correct_j \rightarrow \overline{start}$ whenever $\bar{r}_j(\bar{t}') = r_j(t')$, meaning $(I, r, t') \models K_j (correct_j \rightarrow \overline{start})$. This argument applies to every $j \in G$, hence, (19) is demonstrated for the group G of $2f + 1$ correct agents.

Correctness (C) applied to G at time t' ensures $(I, r, t') \models \bigwedge_{i \in \mathcal{A}} \diamond (correct_i \rightarrow \overline{fire}_i)$, and, since $t \leq t'$, we also have

$$(I, r, t) \models \bigwedge_{i \in \mathcal{A}} \diamond (correct_i \rightarrow \overline{fire}_i).$$

Given that r and t were chosen arbitrarily, we have proved

$$I \models E^{\diamond H} \overline{start} \rightarrow \bigwedge_{i \in \mathcal{A}} \diamond (correct_i \rightarrow \overline{fire}_i).$$

Using (13), we can conclude

$$I \models E^{\diamond H} \overline{start} \rightarrow \bigwedge_{i \in \mathcal{A}} \diamond (correct_i \rightarrow B_i (\overline{start} \wedge E^{\diamond H} \overline{start})),$$

i.e.,

$$I \models E^{\diamond H} \overline{start} \rightarrow \bigwedge_{i \in \mathcal{A}} \diamond H_i (\overline{start} \wedge E^{\diamond H} \overline{start}).$$

In other words, we have demonstrated

$$I \models E^{\diamond H} \overline{start} \rightarrow E^{\diamond H} (\overline{start} \wedge E^{\diamond H} \overline{start}).$$

Using (2) with $\psi = E^{\diamond H} \overline{start}$ and $\phi = \overline{start}$, we conclude

$$I \models E^{\diamond H} \overline{start} \rightarrow C^{\diamond H} \overline{start}. \quad \square$$

³While we only use the fact that agent $j \in G$ is correct at t'_j , these agents will necessarily remain correct throughout run r .