Knowledge-based analysis of the Firing Rebels problem

Krisztina Fruzsa

TU Wien

joint work with Roman Kuznets and Ulrich Schmid

Open Universiteit, Heerlen
November 2, 2021

Motivation

- Firing Rebels with Relay:
  - simplified version of the consistent broadcast primitive [Srikanth and Toueg, JACM87]
  - essentially a non-synchronous version of the Byzantine Firing Squad Problem [Burns and Lynch, 1987]
- Tight connection between knowledge and action in distributed systems:
  - Knowledge of Preconditions Principle [Moses, TARK15]
- Goal: necessary and sufficient knowledge for agents to act

The setting

Our choice:

byzantine fault-tolerant **asynchronous** distributed systems

The setting

Our choice:

byzantine fault-tolerant **asynchronous** distributed systems

Finite set of agents (processing units) $\mathcal{A} = \{1, \ldots, n\}$

- **asynchronous**

The setting

Our choice:

byzantine fault-tolerant **asynchronous** distributed systems

Finite set of agents (processing units) $\mathcal{A} = \{1, \dots, n\}$

- **asynchronous**
- perfect recall

The setting

Our choice:

byzantine fault-tolerant **asynchronous** distributed systems

Finite set of agents (processing units) $\mathcal{A} = \{1, \ldots, n\}$

- **asynchronous**
- perfect recall
- they may be **byzantine** faulty
  - they may deviate from their protocols
  - they may collude to fool other agents
  - false memory

The setting

Our choice:

byzantine fault-tolerant **asynchronous** distributed systems

Finite set of agents (processing units) $\mathcal{A} = \{1, \ldots, n\}$

- **asynchronous**
- perfect recall
- they may be **byzantine** faulty
  - they may deviate from their protocols
  - they may collude to fool other agents
  - false memory

Message-passing communication network

The setting

Our choice:

byzantine fault-tolerant **asynchronous** distributed systems

Finite set of agents (processing units) $\mathcal{A} = \{1, \ldots, n\}$

- **asynchronous**
- perfect recall
- they may be **byzantine** faulty
  - they may deviate from their protocols
  - they may collude to fool other agents
  - false memory

Message-passing communication network

- **asynchronous** (messages can be arbitrarily delayed, i.e., there is no upper bound on message-delivery time)

Firing Rebels with and without Relay

f = maximum number of agents that can turn byzantine in a run

Firing Rebels with and without Relay

f = maximum number of agents that can turn byzantine in a run

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.

Firing Rebels with and without Relay

f = maximum number of agents that can turn byzantine in a run

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.

Example

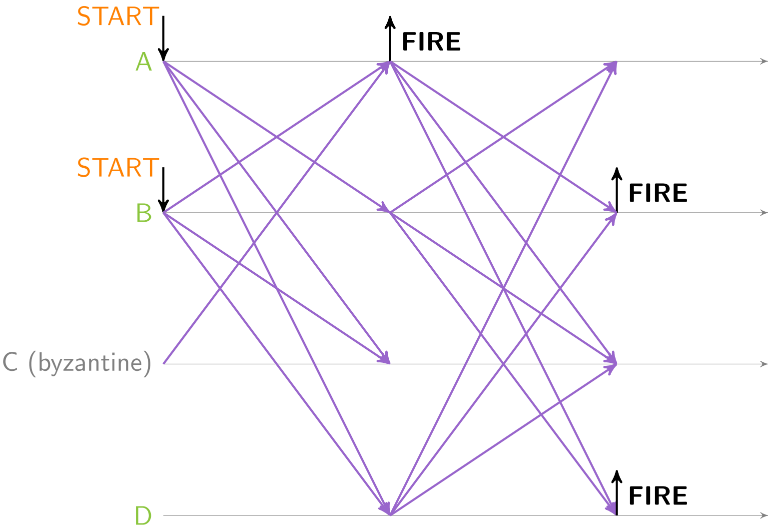

Firing Rebels with and without Relay

f = maximum number of agents that can turn byzantine in a run

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.

The setting

Our choice:

byzantine fault-tolerant asynchronous distributed systems

Preparation for modeling

Consequences of the Brain-in-a-Vat Lemma [Kuznets et al., FroCoS2019]

If at least one agent can become byzantine in a system:

- No agent can ever know that an action or event happened correctly.
- No agent can ever know that it is correct.
- No agent can ever know that another agent is byzantine.

If more than one agent can become byzantine in a system:

- No agent can ever know another agent is correct.

Preparation for modeling

Consequences of the Brain-in-a-Vat Lemma [Kuznets et al., FroCoS2019]

If at least one agent can become byzantine in a system:

- No agent can ever know that an action or event happened correctly.
- No agent can ever know that it is correct.
- No agent can ever know that another agent is byzantine.

If more than one agent can become byzantine in a system:

- No agent can ever know another agent is correct.

Nevertheless, agents can believe.

Runs-and-systems framework

system $=$ set $\mathcal{R}$ of runs

system $=$ set $\mathcal{R}$ of runs

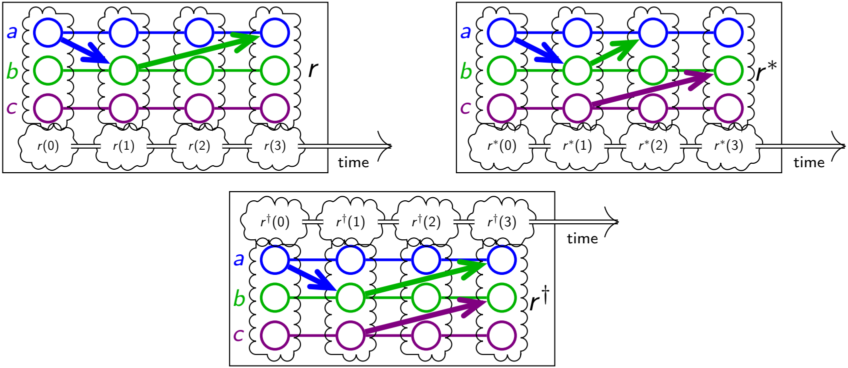

$r(t)$ **global state** at time t in run r

$r_i(t)$ **local state** of agent i at time t in run r

Towards a Kripke model

A point (r, t) refers to time t in run r.
It represents the global state $r(t)$.

Towards a Kripke model

A point (r, t) refers to time t in run r.
It represents the global state $r(t)$.

We want to reason about agents' states of knowledge at various times during a run.

Towards a Kripke model

A point (r, t) refers to time t in run r.
It represents the global state $r(t)$.

We want to reason about agents' states of knowledge at various times during a run.

Therefore:

A point (r, t) is considered a possible world.

Towards a Kripke model

A point (r, t) refers to time t in run r.
It represents the global state $r(t)$.

We want to reason about agents' states of knowledge at various times during a run.

Therefore:

A point (r, t) is considered a possible world.

Two points (r, t) and (r', t') are considered indistinguishable for an agent $i \in \mathcal{A}$ iff $r_i(t) = r'_i(t')$.

Towards a Kripke model

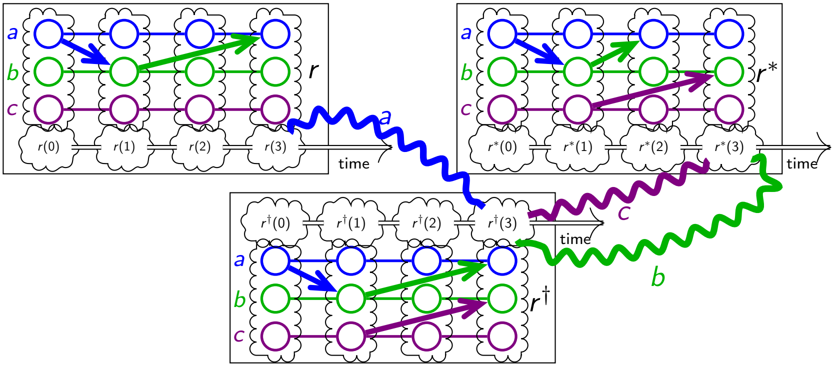

e.g. $r_b^{\dagger}(3) = r_b^{*}(3)$

Syntax

Our language is generated by the following BNF:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_i\varphi \mid \Diamond\varphi \mid Y\varphi,$$

where $p \in Prop$ and $i \in \mathcal{A}$.

For example: $correct_i, \overline{occurred}_i(START) \in Prop$

Syntax

Our language is generated by the following BNF:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_i\varphi \mid \Diamond\varphi \mid Y\varphi,$$

where $p \in Prop$ and $i \in \mathcal{A}$.

For example: $correct_i, \overline{occurred}_i(START) \in Prop$

$$\overline{start}_i := Y\overline{occurred}_i(START) \wedge correct_i$$

$$\overline{start} := \bigvee_{j \in \mathcal{A}} \overline{start}_j$$

$$\overline{fire}_i := \overline{occurred}_i(FIRE) \wedge correct_i$$

$$\overline{fire} := \bigvee_{j \in \mathcal{A}} \overline{fire}_j$$

Syntax

Additional operators we use:

- Belief $B_i\varphi := K_i(correct_i \to \varphi)$ [Moses and Shoham, 1993]
- Hope $H_i\varphi := correct_i \to B_i\varphi$ [F., ESSLLI2019]
- Eventual mutual hope $E^{\Diamond H}\varphi := \bigwedge_{j \in \mathcal{A}} \Diamond H_j\varphi$
- Eventual common hope $C^{\Diamond H}\varphi$ defined as the greatest fixpoint of the equation $\chi \leftrightarrow E^{\Diamond H}(\varphi \wedge \chi)$

Obtaining a Kripke model

A valuation function $\pi : Prop \to 2^{\mathcal{R} \times \mathbb{T}}$ determines at which points $(r, t) \in \mathcal{R} \times \mathbb{T}$ the atomic propositions from *Prop* are *true*.

Obtaining a Kripke model

A valuation function $\pi : Prop \to 2^{\mathcal{R} \times \mathbb{T}}$ determines at which points $(r, t) \in \mathcal{R} \times \mathbb{T}$ the atomic propositions from *Prop* are *true*.

Interpreted system $I = (\mathcal{R}, \pi)$.

Obtaining a Kripke model

A valuation function $\pi : Prop \to 2^{\mathcal{R} \times \mathbb{T}}$ determines at which points $(r, t) \in \mathcal{R} \times \mathbb{T}$ the atomic propositions from *Prop* are *true*.

Interpreted system $I = (\mathcal{R}, \pi)$.

Semantics

- $(I, r, t) \models p$ iff $(r, t) \in \pi(p)$
- $(I, r, t) \models K_i\varphi$ iff $(I, r', t') \models \varphi$ whenever $r'_i(t') = r_i(t)$
- $(I, r, t) \models \Diamond\varphi$ iff $(I, r, t') \models \varphi$ for some $t' \geq t$
- $(I, r, t) \models Y\varphi$ iff $t > 0$ and $(I, r, t - 1) \models \varphi$

Obtaining a Kripke model

A valuation function $\pi : Prop \to 2^{\mathcal{R}\times\mathbb{T}}$ determines at which points $(r, t) \in \mathcal{R} \times \mathbb{T}$ the atomic propositions from *Prop* are *true*.

Interpreted system $I = (\mathcal{R}, \pi)$.

Semantics

- $(I, r, t) \models p$ iff $(r, t) \in \pi(p)$
- $(I, r, t) \models K_i\varphi$ iff $(I, r', t') \models \varphi$ whenever $r'_i(t') = r_i(t)$
- $(I, r, t) \models \Diamond\varphi$ iff $(I, r, t') \models \varphi$ for some $t' \geq t$
- $(I, r, t) \models Y\varphi$ 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 $r \in \mathcal{R}$ and $t \in \mathbb{T}$.

Modeling Firing Rebels

An interpreted system I is consistent with FR for $f \geq 0$ if the following holds:

(C) $I \models \bigvee_{\substack{G \subseteq \mathcal{A} \\ |G|=2f+1}} \bigwedge_{j \in G} B_j \overline{start} \to \bigwedge_{i \in \mathcal{A}} \Diamond(correct_i \to \overline{fire}_i)$

(U) $I \models \overline{fire} \to \overline{start}$

Moreover, I is consistent with FRR if the following holds as well:

(R) $I \models \overline{fire} \to \bigwedge_{i \in \mathcal{A}} \Diamond(correct_i \to \overline{fire}_i)$

Knowledge-based analysis

We wish to know:

- What kind of an epistemic state is **necessary** for a correct agent to be in when firing (for any protocol that meets the requirements of the FR(R) problem specification)?

Necessary state of knowledge

Firing Rebels without Relay

For any interpreted system I consistent with FR and for any agent $i \in \mathcal{A}$:

$$I \models \overline{fire}_i \rightarrow B_i \overline{start}.$$

Necessary state of knowledge

Firing Rebels without Relay

For any interpreted system I consistent with FR and for any agent $i \in \mathcal{A}$:

$$I \models \overline{fire}_i \to B_i \overline{start}.$$

Firing Rebels with Relay

For any interpreted system I consistent with FRR and for any agent $i \in \mathcal{A}$:

$$I \models \overline{fire}_i \to B_i(\overline{start} \land C^{\Diamond H} \overline{start}).$$

Lifting Lemma

Let I be an interpreted system and let $n \geq 3f + 1$.
If I is consistent with FRR, then

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

Knowledge-based analysis

We wish to know:

- What kind of conditions on the interpreted system would be **sufficient** so that the requirements of the FR(R) problem specification are satisfied (i.e., so that the corresponding protocol does meet those requirements)?

Sufficient conditions

For any interpreted system I:

(U) is fulfilled if

$$I \quad \models \quad \bigwedge_{i \in \mathcal{A}} (\neg B_i \overline{start} \rightarrow \neg \overline{fire}_i).$$

Both (U) and (R) are fulfilled if

$$I \models \bigwedge_{i \in \mathcal{A}} \Big((\neg B_i(\overline{start} \wedge C^{\Diamond H} \overline{start}) \rightarrow \neg \overline{fire}_i) \wedge$$

$$\big(B_i(\overline{start} \wedge C^{\Diamond H} \overline{start}) \rightarrow \Diamond(correct_i \rightarrow \overline{fire}_i)\big)\Big).$$

Future work

- Necessary and sufficient communication structures involved in protocols for FR(R)
- Axiomatization of (eventual) common hope

Future work

- Necessary and sufficient communication structures involved in protocols for FR(R)
- Axiomatization of (eventual) common hope

Thank you!