On the Power of Weaker Pairwise Interaction: Fault-Tolerant Simulation of Population Protocols

G. Di Luna*, P. Flocchini*,T. Izumi †,T. Izumi ‡, N. Santoro§, G. Viglietta *
* University of Ottawa, {gdiluna,paola.flocchini,gvigliett}@uottawa.ca
†Nagoya Institute of Technology, t-izumi@nitech.ac.jp
‡Ritsumeikan University, izumi-t@fc.ritsumei.ac.jp
§Carleton University, santoro@scs.carleton.ca

***Abstract*—In this paper we investigate the computational power of population protocols under some unreliable or weaker interaction models. More precisely, we focus on two features related to the power of interactions: omission failures and one-way communications. An omission failure, a notion that this paper introduces for the first time in the context of population protocols, is the loss by one or both parties of the information transmitted in an interaction. The failure may or may not be detected by either party. In one-way models, on the other hand, communication happens only in one direction: only one of the two agents can change its state depending on both agents' states, and the other agent may or may not be aware of the interaction. These notions can be combined, obtaining one-way protocols with (possibly detectable) omission failures.**

We start our investigation by providing a complete classification of all the possible models arising from the aforementioned weaknesses, and establishing the computational hierarchy of these models. We then address for each model the fundamental question of what additional power is necessary and sufficient to completely overcome the model's weakness and make it able to simulate faultless two-way protocols; by "simulator" we mean a wrapper protocol that, implementing an atomic communication of states between two agents, converts any standard two-way protocol into one that works in a weaker model. We answer this question by presenting simulators that work under certain assumptions (e.g., additional memory, unique IDs, etc.) and by proving that simulation is impossible without such assumptions.

I. Introduction

A. Framework

Population protocols [3] are a mathematical model that describes systems of simple mobile computational entities, called *agents*. Two agents can interact (i.e., exchange information) only when their movement brings them into communication range of each other. However, the movements of the agents, and thus the occurrences of their interactions, are completely unpredictable, a condition called "passive mobility". Such would be, for example, the case of a flock of birds, each provided with a sensor; the resulting passively mobile sensor network can then be used for monitoring the activities of the flock and for individual intervention, such as a sensor inoculating the bird with a drug, should a certain condition be detected.

In population protocols, when an interaction occurs, the states of the two agents involved change according to a set of deterministic rules, or "protocol". Interactions are asymmetric: one agent is the "starter", and the other is the "reactor". The execution of the protocol, through the interactions originating from the movements of the entities, generates a non-deterministic sequence of changes in the states of the entities themselves, and thus in the global state of the system. The requirements and goals of a protocol depend on the particular application. In some applications, the goal of the protocol might be to ensure that, in every execution, the system converges to a predefined *final global state*; for example, converging to an "epidemic-alert" state if the number of sensors detecting avian influenza is above a threshold. In other applications, the *sequence of state changes* of each agent must obey precise constraints; for example, if entering some state causes a bird to be inoculated, that state should be entered at most once.

In an interaction, communication is generally assumed to be bidirectional or *two-way*: each agent of a pair receives the state of the other agent and applies the protocol's transition function to update its own state, based on the received information and its current state. From an engineering standpoint, this round-trip communication between two interacting agents may be difficult to implement. Moreover, the standard population protocol model is not resilient to faults.

In this paper we investigate the computational power of population protocols under some unreliable and/or weaker interaction models. More precisely, we focus on two features related to the power of interactions: *omission failures* and *one-way* communications. An omission failure, a notion that this paper introduces for the first time in the context of population protocols, is the loss by one or both parties of the information transmitted in an interaction. The failure may or may not be *detected* by either party. On the other hand, in one-way models (originally introduced in [4]), communication occurs only in one direction: only one of the two agents can change its state depending on both states, and the other agent may or may not be aware of the interaction. These notions can be combined, obtaining one-way protocols with (possibly detectable) omission failures.

A general question is what additional power is necessary and sufficient to fill the gap between the standard two-way model and the weaker models stated above. In this paper we start addressing this question, using as a main investigation tool the concept of a *simulator*: a wrapper protocol converting any protocol for the standard two-way model into one running

under some weaker model. A simulator provides an interface between the simulated protocol and the physical communication layer, giving the system the illusion of being in a two-way environment. As a basic feature, a simulator has to implement an atomic communication of states between two agents, always guaranteeing both safety and liveness of any problem specification; this task is further complicated by the anonymity of the agents, their lack of knowledge of the system, and the limited amount of memory that they may have.

We stress that the importance of focusing on the existence of simulators (as opposed to, say, studying the set of predicates that are computable in a given model) arises from the fact the decisions taken by some agents may be irrevocable by their very nature (e.g., set the value of a write-once register, inoculate a bird, drop a bomb on the target, etc.).

B. Main Contributions

As a first step, we define several omissive models considering both the capability to detect the proximity of another agent (i.e., interacting with another agent without viewing its state) and to detect an omission (i.e, being aware that an interaction failed). Note that these two features are independent. Proximity detection is common in many Medium Access Control protocols, such as CSMA-MPS and STEM.

In the two-way models with omissions, indicated with T_i, detecting the proximity of an agent and not receiving its state is an implicit detection of an omission. For this reason, in these models, we only look at omission detection when: no detection is possible (model T_1), only one side detects it (model T_2), or both sides can detect the omission (model T_3).

In the one-way models, indicated with I_i, detecting the proximity of another agent but not receiving its state could imply either the presence of an omission or that the agent is the starter of the interaction. We consider and analyze all possible combinations of these two capabilities (the complete analysis is included in the Appendix). As a result, we obtain four distinct models: where no detection of omissions is present but the proximity can be detected by the starter (model I_1) or by both parties (model I_2), and where there is detection of proximity by both parties but omissions are detected only by the starter (model I_3) or by the reactor (model I_4). For completeness we also consider the known models without omissions, the original two-way model (TW) and the one-way models introduced in [4]: the *Immediate Transmission* model (IT), where the starter detects the proximity of the reactor, and the *Immediate Observation* model (IO), where there is no detection of proximity.

The hierarchy of these models is shown in Figure 1; more details can be found in Section II.

We consider two main types of *omission adversaries*: a "malignant" one, called **UO**, which can insert omissions at any point in the execution, and a "benign" one, called $\diamond$**NO**, which must eventually stop inserting omissions. Interactions are otherwise "globally fair". Interestingly, all our main simulators work even under the malignant **UO** adversary, while all our main impossibility results hold even under the benign $\diamond$**NO** adversary.

We start by analyzing the negative impact that omissions have on computability. We show that, in the absence of additional assumptions, the simulation of TW protocols in the presence of omissions is impossible even if the agents have infinite memory (Theorem 1). Among other results, we also show that, in the two weak omission models I_1 and I_2, simulation is impossible even under an extremely limited omission adversary, called $\diamond\mathbf{NO_1}$, which can only insert at most one omission in the entire execution.

On the other hand we prove that, in the weakest one-way model, IO, simulation is possible if the agents have unique IDs or the total number of agents, n, is known (Theorems 5 and 6).

In the two strong omission models I_3 and I_4, simulation is possible when an upper bound on the number of omissions is known (Theorem 4). This result in turn implies that, in the non-omissive IT model, TW simulation is possible with a memory overhead of $\Theta(\log n)$ bits for each state of the simulated protocol (Corollary 1). In light of the fact that with constant memory, in absence of additional capabilities, IT protocols are strictly less powerful than two-way protocols [4], our results show that this computational gap can be overcome by using additional memory.

Our main results are summarized in Figure 4, where white blobs represent possibilities and gray blobs impossibilities. As a consequence of these results, we have a complete characterization of the feasibility of simulation when agents have infinite memory, unique IDs, or knowledge of the size of the system. If an upper bound on the number of omissions is known, we also have a complete characterization, except in model T_2, where the problem is still open.

C. Related Work

Since their introduction, there have been extensive investigations on Population Protocols (e.g., see [5], [8], [11], [13], [14], [19]–[21], [26]), and the basic assumptions of the original model have been expanded in several directions, typically to overcome inherent computability restrictions. For example, allowing each agent to have non-constant memory [1], [2], [15]; assuming the presence of a leader [7]; allowing a certain amount of information to be stored on the edges [16]–[18] of the interaction graph.

The issue of dependable computations in population protocols, first raised in [22], has been considered and studied only with respect to processors' faults, and the basic model has necessarily been expanded. In [23] it has been shown how to compute functions tolerating $\mathcal{O}(1)$ crash-stops and transient failures, assuming that the number of failures is bounded and known. In [6] the specific majority problem under $\mathcal{O}(\sqrt{n})$ Byzantine failures, assuming a fair probabilistic scheduler, has been studied. In [27] unique IDs are assumed, and it is shown how to compute functions tolerating a bounded number of Byzantine faults, under the assumption that Byzantine agents cannot forge IDs. Self-stabilizing solutions have been

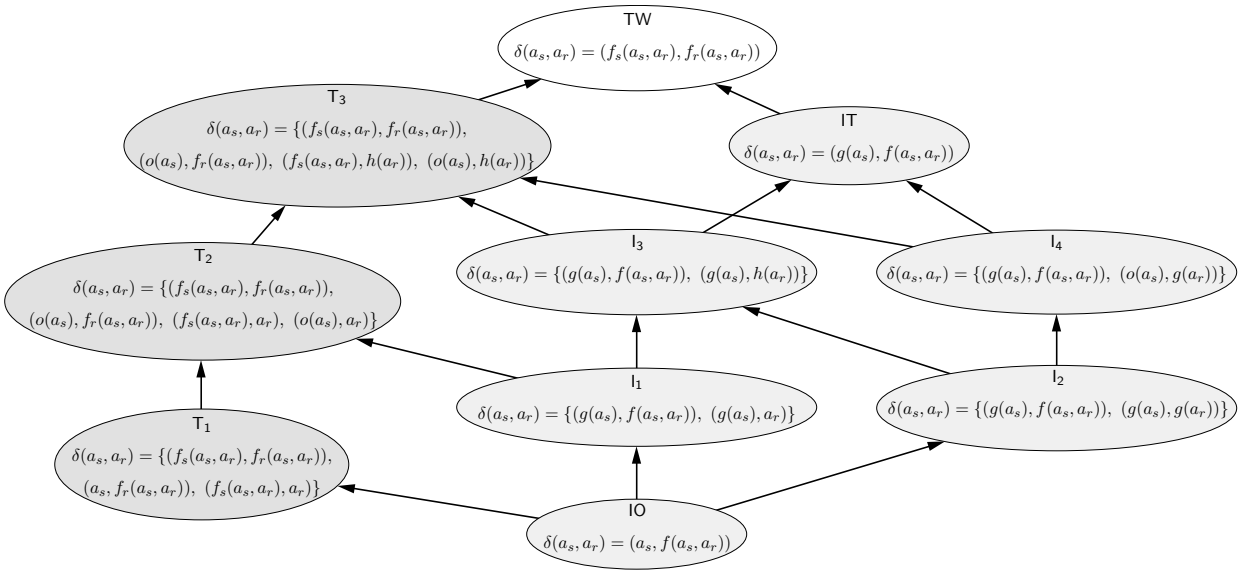

Fig. 1: Computational relationships between models. An arrow between two blobs indicates that the class of problems solvable in the source blob is included in that of the destination blob. The models on the left, T_1, T_2, T_3, are the two-way models with omissions. The models on the right, I_1, I_2, I_3, I_4, are the one-way models with omissions.

devised for specific problems such as leader election (assuming knowledge of the system's size and a non-constant number of states [12], or assuming a leader detection oracle [25]) and counting (assuming the presence of a leader [9]). Moreover, in [10] a self-stabilizing transformer for general protocols has been studied in a slightly different model and under the assumption of unbounded memory and a leader.

Finally, to the best of our knowledge, the one-way model without omissions, has been studied only in [4], where it is shown that IT and IO, when equipped with constant memory, can compute a set of functions that is strictly included in that of TW. Combined with our results in Figure 4, this implies that, without using extra resources (e.g., infinite memory, leader, etc.), simulations are impossible in all the one-way and omissive models.

II. Models and Terminology

A. Population Protocols

We consider a system consisting of a set $A = \{a_1, \ldots, a_n\}$ of mobile agents. The mobility is passive, in the sense that it is decided by an external entity. When two agents meet, they interact with each other and perform some local computation. We always assume that interactions are instantaneous. Each interaction is asymmetric, that is, an interaction between a_s and a_r is indicated by the ordered pair $i = (a_s, a_r)$, where a_s and a_r are called *starter* and *reactor*, respectively. A protocol $\mathcal{P}$ is defined by the following three elements: a set of local states $Q_\mathcal{P}$, a set of initial states $Q'_\mathcal{P} \subseteq Q_\mathcal{P}$, and a transition function $\delta_\mathcal{P} \colon Q_\mathcal{P} \times Q_\mathcal{P} \to Q_\mathcal{P} \times Q_\mathcal{P}$. The function $\delta_\mathcal{P}$ defines the states of the two interacting agents at the end of their local computation. With a small abuse of notation, and when no ambiguity arises, we will use the same literal (e.g., a_i) to indicate both an agent and its internal state. Since the static structure of the system is uniquely determined by $\mathcal{P}$ and n, we refer to it as the *system* $(\mathcal{P}, n)$. A *configuration* C of a system $(\mathcal{P}, n)$ is the n-tuple of local states in $Q_\mathcal{P}$ (i.e., $C \in Q_\mathcal{P}^n$).

Given an k-tuple $t = (x_0, x_1, \ldots, x_{k-1})$ we denote the element x_j by $t[j]$.

Initial Knowledge. To empower the agents, we sometimes assume that each agent has some additional knowledge, such as unique IDs and/or knowledge of n. We model this information by encoding it as a set of initial states of the agents (i.e., in $Q'_\mathcal{P}$).

Executions and Fairness. Whenever an interaction $i = (a_j, a_k)$ turns a configuration of the form $C = (a_1, \ldots, a_j, \ldots, a_k, \ldots, a_n)$ into one of the form

$$C' = (a_1, \ldots, \delta(a_j, a_k)[0], \ldots, \delta(a_j, a_k)[1], \ldots, a_n),$$

we use the notation $C \xrightarrow{i} C'$. A *run* of $\mathcal{P}$ is an infinite sequence of interactions $I = (i_0, i_1, \ldots)$. Given an initial configuration $C_0 \in Q'^n_\mathcal{P}$, each run I induces an infinite sequence of configurations, $\Gamma_I(C_0) = (C_0, C_1, \ldots)$ such that $C_j \xrightarrow{i_j} C_{j+1}$ for every $j \geq 0$, which is called an *execution* of $\mathcal{P}$.

We say that a set of configurations $\mathcal{C} \subseteq Q_\mathcal{P}^n$ is *closed* if, for every $C \in \mathcal{C}$, and for every configuration $\widehat{C}$ obtained by permuting the states of the agents of C, also $\widehat{C} \in \mathcal{C}$.

An execution Γ is *globally fair* (GF) if it satisfies the following condition: for every two (possibly infinite) closed sets of configurations $\mathcal{C}, \mathcal{C}' \subseteq Q_\mathcal{P}^n$ such that for every $C \in \mathcal{C}$ there exists an interaction i and some $C' \in \mathcal{C}'$ such that $C \xrightarrow{i} C'$, if infinitely many configurations of Γ belong to $\mathcal{C}$, then infinitely many configurations of Γ belong to $\mathcal{C}'$ (although not necessarily appearing in Γ as immediate successors of configurations of $\mathcal{C}$).

Note that our definition of global fairness extends the standard one, which only deals with single configurations,

as opposed to sets (see [8]). The two definitions are equivalent when applied to protocols that use only finitely many states, but our extension also works with infinitely many states, while the standard one is ineffective.

B. Interaction Models

In this paper we consider three main models of interactions: the standard *Two-Way* one, and two one-way models presented in [4], i.e., the *Immediate Transmission* model and the *Immediate Observation* model.

Two-Way Model (TW). In this model, any protocol $\mathcal{P}$ must have a state transition function consisting of two functions $f_s \colon Q_\mathcal{P} \times Q_\mathcal{P} \to Q_\mathcal{P}$ and $f_r \colon Q_\mathcal{P} \times Q_\mathcal{P} \to Q_\mathcal{P}$ satisfying $\delta_\mathcal{P}(a_s, a_r) = (f_s(a_s, a_r), f_r(a_s, a_r))$.

Immediate Transmission Model (IT). Any protocol $\mathcal{P}$ must have a state transition function consisting of two functions $g \colon Q_\mathcal{P} \to Q_\mathcal{P}$ and $f \colon Q_\mathcal{P} \times Q_\mathcal{P} \to Q_\mathcal{P}$ satisfying $\delta_\mathcal{P}(a_s, a_r) = (g(a_s), f(a_s, a_r))$ for any $a_s, a_r \in Q_\mathcal{P}$.

Immediate Observation Model (IO). Any protocol $\mathcal{P}$ must have a state transition function of the form $\delta_\mathcal{P}(a_s, a_r) = (a_s, f(a_s, a_r))$.

Note that, in the IT model, the starter explicitly detects the interaction, as it applies function g to its own state. In other terms, even if the starter cannot read the state of the reactor, it can still detect its "proximity". In the IO model, on the other hand, there is no such detection of an interaction (or proximity) by the starter.

C. Omissive Models

An omission is a fault affecting a single interaction. In an omissive interaction an agent does not receive any information about the state of its counterpart. Omissions are introduced by an adversarial entity. We consider:

Definition 1 (Unfair Omissive (UO) Adversary)**.** *The* **UO** *adversary takes a run I and outputs a new sequence I', which is obtained by inserting a (possibly empty) finite sequence of omissive interactions between each pair of consecutive interactions of I.*

Definition 2 (Eventually Non-Omissive ($\diamond$**NO**/$\diamond$**NO$_1$**) Adversary)**.** *The $\diamond$**NO** adversary takes a run I and outputs a new sequence I', which is obtained by inserting any finite sequence of omissive interactions between finitely many pairs of consecutive interactions of I. The $\diamond$**NO$_1$** adversary is even weaker, and can only output interaction sequences with at most one omission.*

If we incorporate omissions in our runs, then transition functions become more general relations.

TW Omissive Model. In the two-way omissive model, we have the transition relation

$$\delta(a_s, a_r) = \{(f_s(a_s, a_r), f_r(a_s, a_r)), (o(a_s), f_r(a_s, a_r)), (f_s(a_s, a_r), h(a_r)),\ (o(a_s), h(a_r))\}$$

(model T_3). The first pair is the outcome of an interaction when no omission is present; the other three pairs represent all possible outcomes when there is an omission: respectively, an omission on the starter's side, on the reactor's side, and on both sides. The functions o and h represent the detection capabilities of each agent: in TW, if one of these is the identity, then omissions are *undetectable* on the respective side.

One-Way Omissive Models. In the case of one-way interactions, we have the transition relation $\delta(a_s, a_r) = \{(g(a_s), f(a_s, a_r)),\ (o(a_s), h(a_r))\}$. The first pair is the outcome of an interaction when no omission is present, and the second pair when there is an omission. (Note that the IO model corresponds to the case in which g is the identity function.) Once again, omissions are undetectable starter-side (respectively, reactor-side) if o (respectively, h) is the identity function.

Hierarchy of Models. The previous models can be weakened by removing the omission detection, either on the starter's side, or on the reactor's side. After identifying all possible combinations of omissions and detections, and pruning out the equivalent ones, the significant models and their relationships have been reported in Figure 1. For TW omissive models, in T_2 we have the models where there is no detection of omission either on the starter's or the reactor's side. Since these two models are symmetric, only the one without reactor-side detection is reported, i.e., function h is forced to be the identity. In T_1 we have the weaker model where no detection is available, i.e., both o and h are the identity. In one-way models, function g is applied when an agent detects the *proximity* of another agent. However, detecting the proximity does not imply the detection of an omission: in I_2, no agent detects an omission, but both detect the proximity of the other agent.

Each arrow in Figure 1 indicates either the obvious inclusion, that is, the transition relation of the source is a special case of the transition relation of the destination, or that the adversary can force the inclusion by avoiding omissions (this is the case with T_3 and TW, for instance). Thus, arrows also indicate inclusions of the sets of problems that are solvable in the various models.

D. Simulation of Two-Way Protocols

In this section we define the *two-way protocol simulator* (or "simulator" for short) and other related concepts. Given a two-way protocol $\mathcal{P}$, consider a protocol $\mathcal{S}(\mathcal{P})$, whose set of local states is $Q_\mathcal{P} \times Q_\mathcal{S}$, where $Q_\mathcal{P}$ is the set of local states of $\mathcal{P}$ (the "simulated states"), and $Q_\mathcal{S}$ is additional memory space used in the simulation. Let $\pi_\mathcal{P} \colon Q_\mathcal{P} \times Q_\mathcal{S} \to Q_\mathcal{P}$ be the projection function onto the set of local states of $\mathcal{P}$. By extension, if C is a configuration of $\mathcal{S}(\mathcal{P})$, we write $\pi_\mathcal{P}(C)$ to indicate the configuration of $\mathcal{P}$ consisting of the projections of the states of the agents of C.

Given an execution $\Gamma_I(C_0)$ of $\mathcal{S}(\mathcal{P})$, where $I = (i_0, i_1, \dots)$, we say that $E(\Gamma) = (e_0, e_1, \dots)$ is a sequence of *events* for Γ if it is a weakly increasing sequence of indices of interactions of I, such that no three indices are the same,

and containing at least the indices of the interactions that determine the update of the simulated state of some agent in the execution Γ (if an interaction updates the simulated states of two agents, then its index must appear twice in $E(\Gamma)$). So, with each event e_j in $E(\Gamma)$, we can associate a unique agent involved in the interaction i_{e_j}; preferably, this agent is one that effectively changes simulated state as a result of i_{e_j}. We also allow extra events in $E(\Gamma)$, associated with agents that do not change simulated state, because we want to take into account simulations of two-way protocols that occasionally leave the state of an agent unchanged.

If $\Gamma_I(C_0) = (C_0, C_1, \dots)$, we let $C_j^- = C_{e_j}$ and $C_j^+ = C_{e_j+1}$. In other words, C_j^- and C_j^+ are the configurations before and after the j-th update of the simulated state, respectively.

Definition 3 (Perfect matching of events)**.** *Given an execution of $\Gamma_I(C_0)$ of a run I and a sequence of events $E(\Gamma)$, a* perfect matching *$M(E)$ is a partition of $\mathbb{N}$ into ordered pairs (viewed as indices of events of $E(\Gamma)$) such that, if $(e_j, e_k) \in M(E)$, where e_j is associated with agent a_x and e_k with agent a_y, then $x \neq y$ and*

$$\delta_{\mathcal{P}}(\pi_{\mathcal{P}}(C_j^-[x]), \pi_{\mathcal{P}}(C_k^-[y])) = (\pi_{\mathcal{P}}(C_j^+[x]), \pi_{\mathcal{P}}(C_k^+[y])).$$

Intuitively, a pair (e_j, e_k) in a perfect matching is the pair of events representing the two state changes given by a two-way interaction of agents under the simulated protocol $\mathcal{P}$. The events e_j and e_k correspond to the updates of the simulated states of the starter and the reactor, respectively. A matching $M(E)$ induces a *derived run* D of $\mathcal{P}$ as follows. Sort the pairs (e_j, e_k) of $M(E)$ by increasing $\min\{e_j, e_k\}$, and let M' be the sorted sequence. Now, if (e'_j, e'_k) is the m-th element of M', agent a_x is associated with event e'_j and agent a_y is associated with event e'_k, then the m-th element of D is (x, y). Now, the *derived execution* induced by $M(E)$ is simply the execution of $\mathcal{P}$ induced by D, i.e., $\Gamma_D(\pi_{\mathcal{P}}(C_0))$.

Definition 4 (Simulation)**.** *A protocol $\mathcal{S}(\mathcal{P})$* simulates *$\mathcal{P}$ if, for any initial configuration C_0 of n agents of $\mathcal{S}(\mathcal{P})$, and any run I whose execution $\Gamma_I(C_0)$ satisfies the* GF *condition, there exists a sequence of events $E(\Gamma)$ with a perfect matching $M(E)$ whose derived execution is an execution of n agents of $\mathcal{P}$ starting from the initial configuration $\pi_{\mathcal{P}}(C_0)$ and satisfying the* GF *condition. We further require that, for each initial configuration C_0, every finite initial sequence of interactions of $\mathcal{S}(\mathcal{P})$ (possibly with omissions) can be extended to an infinite one I, having no additional omissions, whose execution $\Gamma_I(C_0)$ satisfies the* GF *condition.*

The last clause of the definition has been added because, with infinite-memory protocols, the existence of GF executions cannot be taken for granted.

III. Impossibilities for Simulation in Presence of Omissions

In this section, we derive several impossibility results in the presence of omissions. All our impossibility proofs rely on the existence of a two-way protocol that cannot be simulated.

Definition 5 (Pairing Problem)**.** *A set of agents A is given, partitioned into* consumer *agents A_c, starting in state c, and* producer *agents A_p, starting in state p. We say that a protocol $\mathcal{P}$ solves the* Pairing Problem (Pair) *if it enforces the following properties:*

- **Irrevocability.** *$\mathcal{P}$ has a state cs that only agents in state c can get; once an agent has state cs, its state cannot change any more.*
- **Safety.** *At any time, the number of agents in state cs is at most $|A_p|$.*
- **Liveness.** *In all* GF *executions of $\mathcal{P}$, eventually the number of agents in state cs is stably equal to $\min\{|A_c|, |A_p|\}$.*

It is easy to see that Pair can be solved by the simple protocol below in the standard two-way model.

> Pairing Protocol $\mathcal{P}_{IP}$. $Q_{\mathcal{P}_{IP}} = \{cs, c, p, \perp\}$. The only non-trivial transition rules are $(c, p) \mapsto (cs, \perp)$ and $(p, c) \mapsto (\perp, cs)$.

Let us now define a property on the behavior of a generic simulator $\mathcal{S}(\mathcal{P})$ over a sequence of interactions I. We will later show how this property is related to the omission resilience of $\mathcal{S}(\mathcal{P})$.

Definition 6 (Transition Time (TT))**.** *Given a* TW *protocol $\mathcal{P}$, a simulator $\mathcal{S}(\mathcal{P})$, and an execution $\Gamma = (C_0, C_1, \dots)$ of $\mathcal{S}(\mathcal{P})$ on a system of two agents, the* Transition Time (TT) *of the triplet $(\mathcal{S}, \mathcal{P}, \Gamma)$ is the smallest t such that $\pi_{\mathcal{P}}(C_t[0]) = \delta_{\mathcal{P}}(\pi_{\mathcal{P}}(C_0[0]), \pi_{\mathcal{P}}(C_0[1]))[0]$ and $\pi_{\mathcal{P}}(C_t[1]) = \delta_{\mathcal{P}}(\pi_{\mathcal{P}}(C_0[0]), \pi_{\mathcal{P}}(C_0[1]))[1]$ (or ∞, if no such t exists).*

Let $O(I)$ be the number of omissions in a sequence of interactions I.

Definition 7 (Fastest Transition Time (FTT))**.** *Given a* TW *protocol $\mathcal{P}$, a simulator $\mathcal{S}(\mathcal{P})$, and a configuration C_0 for a system of two agents of $\mathcal{S}(\mathcal{P})$, the* Fastest Transition Time (FTT) *of the triplet $(\mathcal{S}, \mathcal{P}, C_0)$ is the smallest* TT *of all the triplets of the form $(\mathcal{S}, \mathcal{P}, \Gamma_I)$, where I ranges over all runs with $O(I) = 0$ and $\Gamma_I[0] = C_0$.*

Intuitively, FTT is the minimum number of (non-omissive) interactions needed by a specific simulator $\mathcal{S}$ to simulate one step of protocol $\mathcal{P}$ in a system of two agents. Thus it can be seen as the "maximum speed" of a simulator. We will show in the following that such a metric is intrinsically related with the omission resilience of $\mathcal{S}$.

A. Impossibilities in Spite of Infinite Memory

In this section we show that simulations of TW models are impossible when omissions are present, even if the system is endowed with infinite memory. We start presenting a key indistinguishability argument.

Lemma 1. *Let $\mathcal{S}(\mathcal{P})$ be a simulator working in the omission model T_3. Let $t > 0$ be the* FTT *of the triplet $(\mathcal{S}, \mathcal{P}, C_0)$, where one agent in C_0 has simulated state q_0, the other agent has q_1 with $q_0 \neq q_1$, and $\delta_{\mathcal{P}}(q_0, q_1) = (q'_0, q'_1)$ and $\delta_{\mathcal{P}}(q_1, q_0) =$*

(q_1', q_0'). *Let A be a system of $2t+2$ agents of $\mathcal{S}(\mathcal{P})$, and let B_0 be an initial configuration of A in which t agents have simulated state q_0 and $t+2$ agents have q_1. Then, there exists a sequence of interactions I^* of A such that $\Gamma_{I^*}(B_0)$ is* GF *and $O(I^*) = t$, with a sequence of events $E(\Gamma_{I^*}(B_0))$ in which at least $t+1$ events represent a transition of some agent from simulated state q_1 to q_1'.*

Proof. Intuitively, we construct a system with t pairs of agents which, thanks to omissive interactions, we "fool" into believing that they are operating in a system of only two agents, until one agent per pair transitions from simulated state q_1 to q_1'. Then we have an extra agent that interacts once with one member of each of the t pairs, also "believing" that the system consists of only two agents, which finally transitions from simulated state q_1 to q_1'. One last auxiliary agent serves as a "generator" of omissive interactions.

Let I be any run of a system of two agents achieving FTT for $(\mathcal{S}, \mathcal{P}, C_0)$; let d_0 be the agent whose initial state is q_0 and let d_1 be the other one. For every $0 \leq k < t$, we construct a sequence of interactions I_k for two agents as follows: copy the first k interactions from I; append an omissive interaction with the same starter as $I[k]$, and with omission (and detection) on d_1's side; extend the resulting sequence to an infinite one whose execution from C_0 satisfies the GF condition, without adding extra omissions (such an extension exists by Definition 4, since $\mathcal{S}$ is a simulator). Note that I_k has exactly one omissive interaction.

Because the execution of I_k is GF, the derived execution must also be GF by definition of simulator, and in particular it makes the simulated states of the two agents transition according to $\delta_{\mathcal{P}}$ infinitely many times. Hence, the agent whose initial simulated state is q_1 will eventually transition to q_1', say after the execution of the first t_k interactions of I_k. Note that this happens regardless of which agent is the starter of the two-way simulated interaction, because by assumption $\delta_{\mathcal{P}}$ is symmetric on (q_0, q_1).

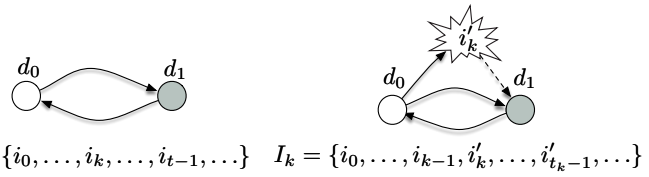

(a) Run I achieving FTT (b) Partial run I_k

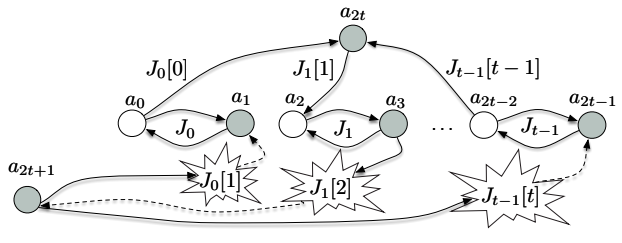

(c) Final run I^*

Fig. 2: Construction of the run I^*

Now name the agents of A as in $A = \{a_0, \ldots, a_{2t+1}\}$, in such a way that, for all $0 \leq k < t$, the agents of the form a_{2k} have simulated state q_0 in B_0, while all other agents of A have q_1. For every $0 \leq k < t$, we construct a sequence J_k, consisting of $t_k + 1$ interactions, involving only agents a_{2k}, a_{2k+1}, a_{2t}, and a_{2t+1}. We make a_{2k} and a_{2k+1} interact with each other as in I_k, but we "redirect" the omissive interaction $I_k[k]$ to a_{2t} and a_{2t+1}. Specifically, we replicate the first k interactions of I_k (where d_0 becomes a_{2k} and d_1 becomes a_{2k+1}); then we add an interaction between a_{2k} and a_{2t}, where the role of a_{2k} (i.e., starter or reactor) is the same as that of d_0 in $I_k[k]$; then we insert an omissive interaction between a_{2k+1} and a_{2t+1}, where the role of a_{2k+1} is the same as that of b_1 in $I_k[k]$, and the omission (and detection) is on a_{2k+1}'s side; finally, we replicate the $t_k - k - 1$ interactions of I_k from $I_k[k+1]$ to $I_k[t_k - 1]$. Observe that J_k contains exactly one omissive interaction, $J_k[k+1]$.

The final sequence I^* is now simply the concatenation of all the sequences J_k (where k goes from 0 to $t-1$, in increasing order), extended to an infinite sequence of interactions whose execution is GF, and having exactly t omissions in total (again, this extension exists by Definition 4).

Let us examine the execution of I^* from the initial configuration B_0. Each of the t pairs of the form (a_{2k}, a_{2k+1}), with $0 \leq k < t$, has an initial execution that is the same as that of (d_0, d_1) interacting for k turns as in I_k. Hence, the execution of a_{2t} is that of d_1 interacting as in I for the first t turns. It follows that a_{2t} transitions from simulated state q_1 to q_1' by the end of the sub-run J_{t-1}. Also, for each pair (a_{2k}, a_{2k+1}), the execution is as in I_k for the first t_k turns; hence, a_{2k+1} transitions from simulated state q_1 to q_1' at the end of the sub-run J_k. Thus, in total, we have at least $t+1$ agents that transition from q_1 to q_1'. □

Theorem 1. *Given an infinite amount of memory on each agent, it is impossible to simulate every* TW *protocol in the* T_3 *model (hence in all the* omissive *models of Figure 1), even under the* $\diamond$**NO** *adversary.*

Proof. We show that the protocol $\mathcal{P}_{IP}$ for Pair cannot be simulated if any type of omissive interaction is allowed. Assume by contradiction that there is a simulator $\mathcal{S}$ for $\mathcal{P}_{IP}$, i.e., $\mathcal{S}$ tsolves Pair under some omissive model. Let us now apply Lemma 1 to $\mathcal{S}$ and $\mathcal{P}_{IP}$, where q_0 is the initial state of the providers (hence there are t providers), q_1 is the initial state of the consumers (hence there are $t+2$ consumers), and q_1' is the irrevocable state.

Because $\mathcal{P}_{IP}$ is symmetric with respect to starter and reactor, the hypotheses of Lemma 1 are satisfied, and hence there is a sequence of interactions I^* whose execution is GF, which causes $t+1$ transitions into the critical state. Since the execution is GF, the derived execution of I^* must be an execution of $\mathcal{P}_{IP}$, due to Definition 4. In particular, it satisfies the irrevocability property of Pair. Therefore, no agent entering a critical state can ever change it. It follows that, eventually, there are at least $t+1$ agents in the critical state, which contradicts the safety property of Pair.

Since I^* contains just finitely many omissive interactions, it can be generated by the $\diamond$**NO** adversary. □

Theorem 1 uses as counterexample the construction of Lemma 1, implying that a simulator S fails to simulate protocol $\mathcal{P}_{IP}$ in a run where the number of failures is exactly the FTT of $(\mathcal{S}, \mathcal{P}_{IP}, (c, p))$. This is even more interesting if we consider simulators that are unaware of the protocol they are simulating, where by "unaware" we mean that the sequence of simulated two-way interactions is not influenced by the protocol that is being simulated or by the initial configuration (i.e., general-purpose and not ad-hoc simulators). We have shown that each of these simulators fails as soon as the number of omissions is above some constant threshold, which is independent of the simulated protocol and the initial configuration. Such a threshold is precisely the minimum number of non-omissive interactions needed to simulate a single two-way transition.

For models T_1,I_1 and I_2, we can strengthen Theorem 1.

Theorem 2. *Given an infinite amount of memory on each agent, it is impossible to simulate every* TW *protocol in the interaction models* T_1*,* I_1*, and* I_2*, even under the* $\diamond$**NO$_1$** *adversary.*

Proof. The proof uses a construction analogous to the one used in Lemma 1. We consider a system $A = \{a_0, \ldots, a_{2t+1}\}$ of $2t + 2$ agents, and we build t sequences of interactions I_k between two agents d_0 and d_1, exactly as in Lemma 1. Recall that the run I_k contains only one omission. Hence, if a simulator is resilient to the $\diamond$**NO$_1$** adversary, it eventually succeeds in making d_0 and d_1 simulate a full two-way interaction, say after t_k one-way interactions. Since t_k is well defined, we can go on and construct the sequence J_k. However, the J_k that we will use in this proof differs from its counterpart used in Lemma 1 by two elements: $J_k[k]$ and $J_k[k+1]$. In particular, our new J_k's will contain no omissions.

If the model is T_1, we replace the old interactions $J_k[k]$ and $J_k[k+1]$ by a single non-omissive interaction between a_k and a_{2t} (in which a_k is the starter if and only if d_0 is the starter in $I[k]$).

Let the model be I_1. If the interaction $I[k]$ is (d_0, d_1), then we replace the old interactions $J_k[k]$ and $J_k[k+1]$ by the single interaction (a_k, a_{2t}). Otherwise, if $I[k] = (d_1, d_0)$, we set $J_k[k] = (a_{2t}, a_{2t+1})$ and $J_k[k+1] = (a_{k+1}, a_{2t+1})$.

Consider now model I_2. If the interaction $I[k]$ is (d_0, d_1), then we set $J_k[k] = (a_k, a_{2t})$ and $J_k[k+1] = (a_{k+1}, a_{2t+1})$. Otherwise, if $I[k] = (d_1, d_0)$, we replace the old interactions $J_k[k]$ and $J_k[k+1]$ by the three interactions (a_{2t}, a_{2t+1}), (a_k, a_{2t+1}), and (a_{k+1}, a_{2t+1}).

Finally, we concatenate the t finite sequences J_k to obtain the new run I^*, which contains no omissions. Let us now examine the execution of I^* from the initial configuration B_0 defined in Lemma 1. Once again, each of the t pairs (a_{2k}, a_{2k+1}), with $0 \leq k < t$, has an initial execution that is the same as that of (d_0, d_1) interacting for k turns as in I_k. Then, the new interactions that we added in lieu of $J_k[k]$ and $J_k[k+1]$ make a_{2k} and a_{2k+1} change state in the same way as in the omissive interaction $I_k[k]$. But as a side effect, also a_{2t} changes state as it would in a non-omissive interaction with a_{2k}. As a consequence, by the end of I^*, all the agents of the form a_{2k+1} with $0 \leq k < t$, as well as a_{2t}, have transitioned from simulated state q_1 to q_1'. Thus, in total, at least $t + 1$ agents transition from q_1 to q_1'.

Now the proof can be completed exactly as in Theorem 1, by showing that the protocol $\mathcal{P}_{IP}$ cannot be simulated. □

One may wonder what would happen if we wanted to construct simulators that "gracefully degrade" when omissions reach a certain threshold t_O. More precisely, for a sequence of interactions I with $O(I) < t_O$, the simulator has to perform a full simulation of $\mathcal{P}$; if $O(I) \geq t_O$, the simulator has to start a simulation, but then it is allowed to stop forever in a "consistent" simulated state. Essentially, in the second case, we allow the sequence of events $E(\Gamma)$ defined in Section II-D to be finite (in other terms, we drop the simulator's "liveness" requirement).

Theorem 3. *Given an infinite amount of memory on each agent, in the* T_3 *model (and hence in all the* omissive *models of Figure 1), any gracefully degrading simulator that simulates all* TW *protocols must have a threshold* $t_O \leq 1$.

Proof. Recall that in Lemma 1 we constructed a sequence of interactions I^* for a set of agents A, which was then applied to the protocol $\mathcal{P}_{IP}$ in order to prove Theorem 1. Suppose now that a simulator has threshold $t_O > 1$. If such a simulator executes a run with at most one omission, it must effectively simulate infinitely many two-way interactions. In particular, it is able to simulate the first two-way interaction in a system of two agents, and therefore the sequence I mentioned in Lemma 1 is well defined for this simulator, as well as the sequences I_k and the numbers t_k. But then, as the agents of A execute the same simulator according to the sequence I^*, they violate the safety property of Pair, because $t + 1$ of them change their simulated state from c to cs. Since they reach a non-consistent simulated state, this means that a gracefully degrading simulator with threshold $t_O > 1$ cannot simulate $\mathcal{P}_{IP}$. □

IV. Simulation in Omissive Models

In this section we focus on designing simulators of two-way protocols. In light of the impossibilities presented in the previous section, additional assumptions are necessary. Section IV-A assumes some knowledge on the maximum number of omissions, Section IV-B assumes the presence of unique IDs, and finally in Section IV-C we assume to know the number of agents.

A. Knowledge on Omissions: Simulator $\mathcal{S}_{KnO}$

Here we assume to know an upper bound o on the number of omissions, i.e., for any sequence of interactions I on which the simulator runs we have $O(I) \leq o$. We will show that under this assumption there exists a simulator for models I_3 and I_4. This contrasts with models I_1 and I_2, in which it is impossible to simulate even when $O(I) \leq 1$ (see Theorem 2).

1: $my\_id = unique\_ID$; $state_{\mathcal{P}} = initial\_state_{\mathcal{P}}$; $id_{other} = \perp$; $state_{other} = \perp$; $state_{sim} = available$ ▷ Agent's variables

2: Upon Event Reactor delivers $(id^s, state^s_{\mathcal{P}}, id^s_{other}, state^s_{other}, state^s_{sim})$
3: **if** $(state_{sim} = available \wedge state^s_{sim} = available)$ **then**
4: $state_{sim} = pairing$
5: $id_{other} = id^s$; $state_{other} = state^s_{\mathcal{P}}$
6: **else if** $(state_{sim} = available \wedge state^s_{sim} = pairing \wedge id^s_{other} = my\_id \wedge state^s_{other} = state_{\mathcal{P}})$ **then**
7: $state_{sim} = locked$
8: $id_{other} = id^s$; $state_{other} = state^s_{\mathcal{P}}$
9: $state_{\mathcal{P}} = \delta_{\mathcal{P}}(state_{\mathcal{P}}, state_{other})[0]$
10: **else if** $(state_{sim} = pairing \wedge id_{other} = id^s \wedge id^s_{other} = my\_id \wedge state^s_{sim} = locked)$ **then**
11: $state_{sim} = available$
12: $id_{other} = state_{other} = \perp$
13: $state_{\mathcal{P}} = \delta_{\mathcal{P}}(state^s_{\mathcal{P}}, state_{\mathcal{P}})[1]$
14: **else if** $(id_{other} = id^s \wedge id^s_{other} \neq my\_id)$ **then**
15: $state_{sim} = available$
16: $id_{other} = state_{other} = \perp$

Fig. 3: Simulation protocol $\mathcal{S}_{ID}$

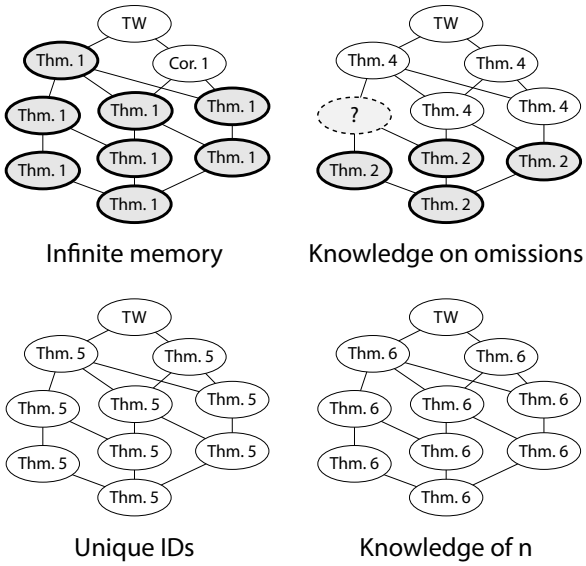

Fig. 4: Map of results (cf. Figure 1)

We explain the simulator $\mathcal{S}_{KnO}$ under model I_3; the version for model I_4 is only slightly different, and its correctness follows from symmetry considerations. The simulator is based on the exchange of "tokens". Each simulated state $q \in Q_{\mathcal{P}}$ is represented as a sequence of numbered tokens: $\langle q, 1\rangle, \ldots, \langle q, o+1\rangle$. Intuitively, an agent tries to transmit its state to others by sending one token at a time, for $o+1$ consecutive interactions, each time incrementing the counter. When a reactor detects an omission, it generates a *joker* token $\langle J\rangle$, which will also be sent in successive interactions. Note that there are never more than o jokers circulating. Every time an agent gets a new token, it checks if it owns the complete set of $o+1$ tokens representing some state q and, if so, it simulates (part of) an interaction with a hypothetical partner in state q. If the complete set of tokens is not available, the agent is allowed to replace the missing tokens with the jokers that it currently owns. After the $o+1$ tokens have been used, they are discarded and withdrawn from circulation. However, if an agent uses some joker tokens, it "takes note" of what tokens these jokers are replacing. If later on the same agent obtains one of the tokens in this list, say $\langle q, i\rangle$, it turns $\langle q, i\rangle$ into a joker and removes $\langle q, i\rangle$ from the list. (This is reminiscent of the card game Rummy.)

Simulator Variables. Each agent has a queue of tokens to be sent, called $sending$, initially empty. It also has a variable $state_{sim} = \mathsf{available}$ for the state of the simulator protocol, a variable $state_{\mathcal{P}}$ for the state of the simulated protocol (initialized according to its initial simulated state), and a multi-set of tokens called $Jokers$, initially empty.

Simulator Protocol. Suppose that an agent interacts as a starter. If $state_{sim} = \mathsf{available}$ and $sending$ is empty, the agent switches to $state_{sim} = \mathsf{pending}$ and inserts the complete set of tokens $\langle state_{\mathcal{P}}, 1\rangle, \ldots, \langle state_{\mathcal{P}}, o+1\rangle$ into $sending$. In any case, and regardless of $state_{sim}$, the starter removes the first token from the queue, and the reactor reads it.

Suppose now that an agent interacts as a reactor. To begin with, it reads the first token from the $sending$ queue of the starter, and enqueues it into its own $sending$ queue. If it detects an omission, it enqueues a joker token instead. Then it performs a preliminary check: if $state_{sim} = \mathsf{pending}$, and the agent can find a complete set of tokens for its own state (i.e., $state_{\mathcal{P}}$) in its own $sending$ queue (possibly using some joker tokens as wildcards), it switches to $state_{sim} = \mathsf{available}$ and removes the set of $o+1$ used tokens from the queue. After this preliminary check, the core protocol starts: if $state_{sim} = \mathsf{available}$, and the agent has a complete set of tokens for some state q in its own $sending$ queue (possibly using some joker tokens), it removes the set of $o+1$ used tokens from the queue, it simulates its part of the two-way transition with an agent in state q (i.e., it updates $state_{\mathcal{P}} = \delta(q, state_{\mathcal{P}})[1]$), and it enqueues into $sending$ a complete set of "state change" tokens, i.e., $\langle (q, state_{\mathcal{P}}), 1\rangle, \ldots, \langle (q, state_{\mathcal{P}}), o+1\rangle$. On the other hand, if $state_{sim} = \mathsf{pending}$, and the agent has a complete set of state change tokens of the form $\langle (state_{\mathcal{P}}, q'), i\rangle$ in its own $sending$ queue (possibly using some joker tokens), it removes the set of $o+1$ used tokens from the queue, it updates $state_{\mathcal{P}} = \delta(state_{\mathcal{P}}, q')[0]$, and switches to $state_{sim} = \mathsf{available}$.

Also, whenever a reactor uses a joker token as a substitute for some token $\langle q, i\rangle$, it adds $\langle q, i\rangle$ to the multi-set $Jokers$. Symmetrically, when it receives a new token $\langle q, i\rangle$ from a starter and that token is in $Jokers$, it removes one copy of $\langle q, i\rangle$ from $Jokers$, removes the last copy of $\langle q, i\rangle$ from $senders$, and enqueues a new joker token into $senders$.

Due to space constraints, the proof of correctness of this simulator is in the Appendix.

Theorem 4. *Given an upper bound o on the number of omissions and $\Theta(|Q_{\mathcal{P}}|(o+1)\log n)$ bits of memory on each agent, every* TW *protocol can be simulated in* I_3 *and* I_4. □

By applying this theorem to a system without omissions (i.e., plugging $o = 0$), we have:

Corollary 1. *Given $\Theta(|Q_{\mathcal{P}}| \log n)$ bits of memory on each agent, every* TW *protocol can be simulated in* IT. □

B. Unique IDs and IO*: Simulator $\mathcal{S}_{ID}$.*

Now we assume that the agents have unique IDs as part of their initial state, and we give a TW simulator for the IO model, named $\mathcal{S}_{ID}$, which is reported in Figure 3. The idea is to use the uniqueness of the IDs to implement a locking mechanism that ensures the consistent matching of simulated state changes. Essentially, at a certain point an agent commits itself to executing a transition only with another agent with a specific ID. The locking scheme contains a rollback procedure to avoid deadlocks.

Simulator Variables. Each agent has the following variables: my_id for its own ID, $state_{sim} =$ avaible for the state of the simulator protocol, and $state_{\mathcal{P}}$ for the state of the simulated protocol. Moreover, it keeps two variables, id_{other} and $state_{other}$, which are the ID and the state of the other agent in the simulated two-way interaction.

Simulator Protocol. When an available reactor a_r, with ID r and simulated state $state_{\mathcal{P}} = q_r$, observes a starter a_s with ID s and $state^s_{sim} =$ available, it enters a pairing state. Moreover, it saves the ID s in id_{other} and the simulated state $state^s_{\mathcal{P}} = q_s$ of a_s in $state_{other}$ (see the details at Lines 3–5). The pairing state could be seen as a "soft" commitment in which a reactor picks a specific agent as a possible partner for a two-way interaction. In some specific conditions, an agent in the pairing state can "roll back" to the available state without completing a simulated two-way interaction; this will be covered later.

The simulation proceeds as soon as a_s, which is available, receives the information that some other agent a_r is in the pairing state and wants to pair up with an agent that has $my\_id = s$ and simulated state q_s. In this case a_s sets its simulator state to locked, stores a_r's simulated state and ID, and executes the transition $\delta_{\mathcal{P}}(state_{\mathcal{P}}, state_{other} = q_r)[0] = f_s(state_{\mathcal{P}}, state_{other})$. We remark that this happens only if the current simulated state of a_s is equal to the variable $state_{other}$ of a_r (see Line 6).

Suppose that a_s is locked; if a_r observes a_s, it executes the transition $\delta_{\mathcal{P}}(state^s_{\mathcal{P}}, state_{\mathcal{P}})[1] = f_r(state^s_{\mathcal{P}}, state_{\mathcal{P}})$, becomes available, and resets the variable id_{other} (see Lines 10–13). Now, if a_s is locked and observes that a_r's variable id_{other} is not s, then it resets its own state to available (see Lines 14–16).

It may happen, due to the IO model's nature, that a_s, with variable $state_{\mathcal{P}} = q_s$, induces an agent a_r to enter state pairing, but then a_s starts a two-way simulation with a different agent. In order to prevent a_r from waiting forever, we make it reset the pending transition if it encounters a_s again with $id^s_{other} \neq my\_id$ (this is incorporated in Lines 14–16).

Theorem 5. *Assuming* IO *and unique IDs, $\mathcal{S}_{ID}$ is a* TW *simulator.*

Proof. Let us consider the simulation of a generic two-way protocol $\mathcal{P}$. Assume that an agent a_0 becomes pairing upon observing an agent a_1. Later, a_1 can either become locked with a_0, or pairing as well, upon observing some other agent a_2. It is clear that, if such a "chain" of pairing agents is formed, it must stop eventually. The last agent in the chain, say a_k, will then have to become locked upon observing some pairing agent with id_{other} equal to a_k's ID (which will eventually happen due to the GF condition).

Now, whenever an agent a_s enters state locked after observing an agent a_r in state pairing, it changes its simulated state according to $\delta_{\mathcal{P}}$, say at time t_s, and sooner or later also a_r will do the same, say at time t_r, with $t_r > t_s$. This is because a_r cannot start a new interaction with a_s between times t_s and t_r (since a_s would have to be in state available), and hence it will necessarily be seen by a_s with $id_{other} \neq s$, due to the GF condition. Moreover, a_s cannot change its own simulated state after t_s and before t_r, because it is locked.

We have proved that infinitely many simulated state transitions must occur; these events can easily be paired up into a consistent perfect matching. We only have to prove that the derived execution satisfies the GF condition. We will do it in the case in which the system consists of $n \geq 3$ agents; the proof for the case $n = 2$ is simpler, and we omit it. Let $\mathcal{C}$ and $\mathcal{C}'$ be closed sets of configurations of $\mathcal{P}$, such that every configuration of $\mathcal{C}$ can become one of $\mathcal{C}'$ after a two-way interaction, and suppose that the derived execution passes through $\mathcal{C}$ infinitely many times. Let $\widetilde{\mathcal{C}}$ be the set of configurations of the simulator protocol whose simulated states are in $\mathcal{C}$ and let $\widetilde{\mathcal{C}}'$ be constructed similarly from $\mathcal{C}'$. (Note: if a configuration of the simulator protocol contains a locked agent a_s, the simulated state of its partner a_r is assumed to be the state it would reach after the interaction with a_s. This agrees with the definition of derived run given in Section II-D.) By assumption, the simulation passes through $\widetilde{\mathcal{C}}$ infinitely often; we claim that it must go through $\widetilde{\mathcal{C}}'$ infinitely many times, as well. By definition of $\mathcal{C}$, for every $C_j \in \widetilde{\mathcal{C}}$, there is an interaction in $\mathcal{P}$ between two agents a_s and a_r that maps $\pi_{\mathcal{P}}(C_j)$ into $\pi_{\mathcal{P}}(C'_j)$, where $C'_j \in \widetilde{\mathcal{C}}'$. We will prove that such a C'_j can be reached from C_j after at most a constant number of interactions.

- If a_s is available in C_j and a_r is either available or pairing with a_s, then C'_j can be obtained by simply letting a_s and a_r interact together multiple times until they perform a full simulated interaction, and their states transition according to $\delta_{\mathcal{P}}$.
- If a_s or a_r (perhaps both) is locked in C_j, we let it interact with its current partner until the simulated interaction is completed and its internal state is again available. Then we proceed as in the other cases.
- If a_s is pairing in C_j or a_r is pairing with an agent that is not a_s, we have to make it become available without performing a full two-way interaction, and then we can proceed as in the other cases. Suppose that a_s is pairing (the case with a_r is handled similarly), and let a_q be the agent with which a_s is paired (perhaps $a_q = a_r$).

- If a_q is pairing in C_j (of course not with a_s), then we let a_s observe a_q and roll back to the available state.
- If a_q is available in C_j, we let it pair up with some other available agent (possibly a_r), and the we proceed as in the previous case. If an available agent does not exist, we can create one by letting some pairing agent roll back or some locked agent complete its current interaction, as explained in the first paragraph of the proof.
- If a_q is locked in C_j, we let it finish the simulated interaction and become available. If it was locked with a_s, we are finished because now a_s is available too. Otherwise, we proceed as in the previous case.

As already observed, C'_j can be reached from C_j after at most a constant number c of interactions, and this holds for every j. By applying the definition of GF to the simulator's execution c times, we have that C'_j is indeed reached for infinitely many j's. Therefore $\widetilde{\mathcal{C}'}$ is reached infinitely many times, and thus so is $\mathcal{C}'$ by the derived execution. □

C. Simulating with knowledge of n

We give the following additional result on simulating when additional knowledge is available to the agents. The proof uses a naming algorithm in conjunction with $\mathcal{S}_{ID}$. Due to lack of space, the details can be found in the Appendix.

Theorem 6. *With the knowledge of $|A| = n$ and $\Theta(\log n)$ bits of memory on each agent, every* TW *protocol can be simulated in* IO. □

V. Conclusion

In this paper we have given a formal definition of two-way simulation in population protocols, and we identified several omission models. On top of this framework, we have given several impossibility results, as well as two-way simulators. Our results yield an almost comprehensive characterization, see Figure 4. The only gap left concerns the possibility of simulation in model T_2 when an upper bound on the number of omissions is known. As future work we are going to investigate this gap and study models where a unique leader agent is present. Our preliminary results, [24], in the latter direction show that the problem is far from trivial, and two-way simulation is still impossible in a wide set of models.

References

[1] D. Alistarh and R. Gelashvili. "Polylogarithmic-time leader election in population protocols", *42th International Colloquium on Automata, Languages and Programming, ICALP*, 2015, pp. 479–491.

[2] D. Alistarh, R. Gelashvili, and M. Vojnovic. "Fast and exact majority in population protocols", *34th Annual ACM Symposium on Principles of Distributed Computing, PODC*, 2015, pp. 47–56.

[3] D. Angluin, J. Aspnes, Z. Diamadi, M. J. Fischer, and R. Peralta. "Computation in networks of passively mobile finite-state sensors", *Distributed Computing*, Vol. 18(4), 2006, pp. 235–253.

[4] D. Angluin, J. Aspnes, and D. Eisenstat. "On the power of anonymous one-way communication", *9th International Conference on Principles of Distributed Systems, OPODIS*, 2005, pp. 396–411.

[5] D. Angluin, J. Aspnes, and D. Eisenstat. "Stably computable predicates are semilinear", *25th Annual ACM Symposium on Principles of Distributed Computing, PODC*, 2006, pp. 292–299.

[6] D. Angluin, J. Aspnes, and D. Eisenstat. "A simple population protocol for fast robust approximate majority", *Distributed Computing*, Vol. 21(2), 2008, pp. 87–102.

[7] D. Angluin, J. Aspnes, and D. Eisenstat. "Fast computation by population protocols with a leader", *Distributed Computing*, Vol. 21(3), 2008, pp. 61–75.

[8] J. Aspnes and E. Ruppert. "An introduction to population protocols", *Bulletin of the European Association for Theoretical Computer Science*, Vol. 93, 2007, pp. 98–117.

[9] J. Beauquier, J. Burman, S. Clavière, and D. Sohier. "Space-optimal counting in population protocols", *29th International Symposium on Distributed Computing, DISC*, 2015, pp. 631–649.

[10] J. Beauquier, J. Burman, and S. Kutten. "A self-stabilizing transformer for population protocols with covering", *Theoretical Computer Science*, Vol. 412(33), 2011, pp. 4247–4259.

[11] O. Bournez, P. Chassaing, J. Cohen, L. Gerin, and X. Koegler. "On the convergence of population protocols when population goes to infinity", *Applied Mathematics and Computation*, Vol. 215(4), 2009, pp. 1340–1350.

[12] S. Cai, T. Izumi, and K. Wada. "How to prove impossibility under global fairness: on space complexity of self-stabilizing leader election on a population protocol model", *Theory of Computing Systems*, Vol. 50(3), 2012, pp. 433–445.

[13] D. Canepa, and M. Gradinariu Potop-Butucaru. "Self-stabilizing tiny interaction protocols", *3rd International Workshop on Reliability, Availability, and Security, WRAS*, pp. 1–6, 2010.

[14] I. Chatzigiannakis, S. Dolev, S. P. Fekete, O. Michail, and P. G. Spirakis. "Not all fair probabilistic schedulers are equivalent", *13th International Conference on Principles of Distributed Systems, OPODIS*, 2009, pp. 33–47.

[15] I. Chatzigiannakis, O. Michail, S. Nikolaou, and A. Pavlogiannis. "Passively mobile communicating machines that use restricted space", *Theoretical Computer Science*, Vol. 412(46), 2011, pp. 6469–6483.

[16] I. Chatzigiannakis, O. Michail, S. Nikolaou, A. Pavlogiannis, and P. G. Spirakis. "All symmetric predicates in NSPACE(n^2) are stably computable by the mediated population protocol model", *35th International Symposium on Mathematical Foundations of Computer Science, MFCS*, 2010, pp. 270–281.

[17] I. Chatzigiannakis, O. Michail, and P. G. Spirakis. "Stably decidable graph languages by mediated population protocols", *12th International Symposium on Stabilization, Safety, and Security of Distributed Systems, SSS*, 2010, pp. 252–266.

[18] I. Chatzigiannakis, O. Michail, and P. G. Spirakis. "Mediated population protocols", *Theoretical Computer Science*, Vol. 412(22), 2011, pp. 2434–2450.

[19] I. Chatzigiannakis, O. Michail, and P. G. Spirakis. "New models for population protocols", *Synthesis Lectures on Distributed Computing Theory*, Morgan & Claypool, 2011.

[20] I. Chatzigiannakis and P. G. Spirakis. "The dynamics of probabilistic population protocols", *22th International symposium on Distributed Computing, DISC*, 2008, pp. 498–499.

[21] H.-L. Chen, R. Cummings, D. Doty, and D. Soloveichik. "Speed faults in computation by chemical reaction networks", *28th International Symposium on Distributed Computing, DISC*, 2014, pp. 16–30.

[22] C. Delporte-Gallet, H. Fauconnier, and R. Guerraoui. "What dependability for networks of mobile sensors?" *In Proceedings of the First Workshop on Hot Topics in System Dependability, HotDep*, 2005, p. 8.

[23] C. Delporte-Gallet, H. Fauconnier, R. Guerraoui, and E. Ruppert. "When birds die: making population protocols fault-tolerant", *2nd IEEE International Conference on Distributed Computing in Sensor Systems, DCOSS*, 2006, pp. 51–66.

[24] G. A. Di Luna, P. Flocchini, T. Izumi, T. Izumi, N. Santoro, and G. Viglietta. "Population protocols with faulty interactions: the impact of a leader", *arXiv:1611.06864 [cs.DC]*, 2016.

[25] M. Fischer and H. Jiang. "Self-stabilizing leader election in networks of finite-state anonymous agents", *10th International Conference on Principles of Distributed Systems, OPODIS*, 2006, pp. 395–409.

[26] R. Guerraoui and E. Ruppert. "Even small birds are unique: population protocols with identifiers", *Technical Report CSE-2007-04, York University*, 2007.

[27] R. Guerraoui and E. Ruppert. "Names trump malice: tiny mobile agents can tolerate byzantine failures", *36th International Colloquium on Automata, Languages and Programming, ICALP*, Vol. 16(2), 2009, pp. 484–495.

APPENDIX

Model Investigation for One-Way Interactions

All possible one-way interaction models are detailed in Figure 5. The figure contains all possible combinations of detection of proximity and starter- or reactor-side omissions. In the following we will show the immediate equivalences between these models, which will result in only the four significant one-way omissive models contained in Figure 1.

Consider a transition relation in the form $\delta(a_s, a_r) = \{(g(a_s), f(a_s, a_r)), (o(a_s), h(a_r)/a_r)\}$. In this case agents simulate a model without omissions by using the identity function for h, o. Thus, all these models are equivalent to IT. Since omissions are introduced by an adversary, it is clear that the adversary avoids inserting omissions when these give more computational power to the agents, i.e., if they break the symmetry. Therefore, all models such that $\delta(a_s, a_r) = \{(a_s, f(a_s, a_r)), (o(a_s), (g/h(a_r))/a_r)\}$ are equivalent to IO.

Detection Capabilities	Transition Relation
Str(PD)	$\delta(a_s, a_r) = \{(g(a_s), f(a_s, a_r)), (g(a_s), a_r)\}$
Str(OD)	$\delta(a_s, a_r) = \{(a_s, f(a_s, a_r)), (o(a_s), a_r)\}$
Str(PD+OD)	$\delta(a_s, a_r) = \{(g(a_s), f(a_s, a_r)), (o(a_s), a_r)\}$
Rct(PD)	$\delta(a_s, a_r) = \{(a_s, f(a_s, a_r)), (a_s, g(a_r))\}$
Rct(OD)	$\delta(a_s, a_r) = \{(a_s, f(a_s, a_r)), (a_s, h(a_r))\}$
Rct(PD+OD)	$\delta(a_s, a_r) = \{(a_s, f(a, a_r)), (a_s, h(a_r))\}$
Str(PD) & Rct(PD)	$\delta(a_s, a_r) = \{(g(a_s), f(a_s, a_r)), (g(a_s), g(a_r))\}$
Str(PD) & Rct(OD)	$\delta(a_s, a_r) = \{(g(a_s), f(a_s, a_r)), (g(a_s), h(a_r))\}$
Str(PD) & Rct(PD+OD)	$\delta(a_s, a_r) = \{(g(a_s), f(a_s, a_r)), (g(a_s), h(a_r))\}$
Str(OD) & Rct(PD)	$\delta(a_s, a_r) = \{(a_s, f(a_s, a_r)), (o(a_s), g(a_r))\}$
Str(OD) & Rct(OD)	$\delta(a_s, a_r) = \{(a_s, f(a_s, a_r)), (o(a_s), h(a_r))\}$
Str(OD) & Rct(PD+OD)	$\delta(a_s, a_r) = \{(a_s, f(a_s, a_r)), (o(a_s), h(a_r))\}$
Str(PD+OD) & Rct(PD)	$\delta(a_s, a_r) = \{(g(a_s), f(a_s, a_r)), (o(a_s), g(a_r))\}$
Str(PD+OD) & Rct(OD)	$\delta(a_s, a_r) = \{(g(a_s), f(a_s, a_r)), (o(a_s), h(a_r))\}$
Str(PD+OD) & Rct(PD+OD)	$\delta(a_s, a_r) = \{(g(a_s), f(a_s, a_r)), (o(a_s), h(a_r))\}$

Fig. 5: Various models of one-way interactions in the presence of proximity detection (PD) and omission detection (OD) for Starter (Str) and Reactor (Rct).

Proof of Theorem 4

Lemma 2. *The derived execution of simulator $\mathcal{S}_{KnO}$ is* GF.

Proof. Let $\mathcal{C}$ and $\mathcal{C}'$ be closed sets of configurations of $\mathcal{P}$, such that every configuration of $\mathcal{C}$ can become one of $\mathcal{C}'$ after a two-way interaction, and suppose that the derived execution passes through $\mathcal{C}$ infinitely many times. Let $\widetilde{\mathcal{C}}$ be the set of configurations of the simulator protocol whose simulated states are in $\mathcal{C}$, and let $\widetilde{\mathcal{C}}'$ be constructed similarly from $\mathcal{C}'$. By assumption, the simulation passes through $\widetilde{\mathcal{C}}$ infinitely often; we claim that it must go through $\widetilde{\mathcal{C}}'$ infinitely many times, as well. By definition of $\mathcal{C}$, for every $C_j \in \widetilde{\mathcal{C}}$, there is an interaction in $\mathcal{P}$ between a_s and a_r that maps $\pi_{\mathcal{P}}(C_j)$ into $\pi_{\mathcal{P}}(C'_j)$, where $C'_j \in \widetilde{\mathcal{C}}'$ is obtained from C_j by changing the simulated states of a_s and a_r according to $\delta_{\mathcal{P}}$ (and possibly some simulator variables). Since in C_j there are no "pending" transactions, C_j has a possible continuation $\widehat{C}_j$ where agents a_s and a_r are available, such that $\pi_{\mathcal{P}}(\widehat{C}_j) = \pi_{\mathcal{P}}(C_j)$, and from this configuration the final configuration $\widetilde{\mathcal{C}}'$ is reachable. The proof is done by case analysis on the number of agents:

- Case $n > 4$: In this case a simple counting arguments shows that it is always possible to find a finite sequence of interactions such that: agents a_s and a_r are available and have an empty buffer and no agent sets $state_{\mathcal{P}} = \delta(q_s, state_{\mathcal{P}})[1]$. Let q_r, q_s be the state of a_r, a_s, respectively. Having at most $n(o+1)$ tokens, distributing this tokens among $n-2$ agents, such that each agent obtains at most $(1+\frac{2}{n-2})(o+1) < (1+2/3)(o+1)$ tokens. No agent can set $state_{\mathcal{P}} = \delta(q_s, state_{\mathcal{P}})[1]$ if it does not receives at least $2(o+1)$ tokens, please note that by assumption we started by a configuration where agents were either $pending$ or $available$ and no run $< (q_s, *), * >$ was present. Starting from this configuration, we first show how to "unlock" agent a_r (or a_s if it is pending). If agent a_r is pending there is an agent a_x containing the first token of the run $< q_r, * >$, this token can be in position j in the buffer of a_x with $0 < j < (1 + 2/3)(o + 1)$. If $j = 0$ the token is sent to a_r by doing an interaction (a_x, a_r) otherwise we do first a sequence of $j - 1$ interactions (a_x, a_r), note that this could at most trigger the exit of a_r from the $pending$ state since a_r receives strictly less then $2(o + 1)$ tokens, if this is the case we have done. In any case we append another sequence of interactions (a_r, a_x) until the buffer of a_r is empty. Now either a_r is available, or the token in position $j = 0$ in the buffer of a_x is the first token of the run of a_r, in this case (a_x, a_s) moves this token to the buffer of a_s. This procedure is iterated until all tokens but one are in the buffer of a_s, the last token is in position $j = 0$ of agent a_y and the buffer of a_r is empty, when this happens a trivial sequence of interactions (a_s, a_r) followed by (a_y, a_r) bring a_r in state $available$ with an empty buffer and it empties the buffer of a_s. The same procedure can be used to unlock a_s if it is pending. Once a_s and a_r are both available and with an empty buffer a sequence of $(o + 1)$ interactions (a_s, a_r) followed by a sequence of $(o+1)$ interactions (a_r, a_s) brings the system in configuration $\widetilde{\mathcal{C}}'$.
- Case $n = 4, 3$: Let us first examine the case of three agents, and let a_x with state q_x be the third agent. Let us assume that a_r does not have in the buffer a complete run for $< q_x, * >$ otherwise it sends its tokens to a_x until the first token of the run for q_x is sent to a_x. Then we send all tokens for the run $< q_r, * >$ to a_r: if agent a_x has some token for $< q_s, * >$ or $< q_r, * >$ in position 0 of the buffer we have an interaction (a_x, a_r), the same is done for agent a_s this procedure stop when both agents a_s, a_x have a token for $< q_x, * >$ in position 0. When this happen we let agents have interactions (a_s, a_x) and (a_s, a_x), this allows the agent to remove the tokens for q_x, note that during this procedure agent a_s does not increases its number of tokens for $< q_x, * >$ contained in buffer. Iterating this procedure agent a_r obtains a complete run $< q_r, * >$ and $< q_s, * >$. At this points it exits from

pending being available and it executes the first portion of the simulated two interactions. A series of successive $o+1$ interactions (a_r, a_s) brings the system in configuration $\widetilde{\mathcal{C}}'$. The case for $n = 4$ agents is analogous.

- Case $n = 2$: In this case after at most $2(o+1)$ interactions (a_s, a_r) and after other $(o+1)$ interactions (a_r, a_s) we have that the agents changes simulated state bringing the system in configuration $\widetilde{\mathcal{C}}'$.

□

Theorem 4. *Assuming* I_3 *or* I_4 *and* $\Theta(\log n|Q_{\mathcal{P}}|(o+1))$ *bits of memory on each agent, there exists a protocol that simulates every* TW *protocol.*

Proof. The proof uses model I_3, the correctness for model I_4 follows from symmetry consideration. We first show that an agents sets $state_{\mathcal{P}} = \delta(q, state_{\mathcal{P}})[1]$ infinitely many times. Let us first consider the case where an agent a_r with $state_{sim} =$ available exists, if there exists an agent $a_s \neq a_r$ with state $state_{sim} =$ available it is easy to see that after $o+1$ interactions (a_s, a_r) , despite the presence of omissions, agent a_r will have a run for the state q_s of a_s in $sending$, therefore it executes $state_{\mathcal{P}} = \delta(q_s, state_{\mathcal{P}})[1]$. In case a_s is in pending and there is no token $< (q_s, *), * >$, we show that there exists a run for a state a_s scattered among agents.

We claim that once a run is created it disappears from the system only if it is consumed by an agent. Suppose that there is no token $< q_s, * >$ this implies that each time an agent was trying to transmit a toke for state q_s an omission occurred, but there are $o+1$ such tokens and at most o omissions, therefore this is impossible. Now let us suppose that the number of jokers and tokens $< q_s, * >$ is less than $o+1$: each time an omission occurred when a token $< q_s, * >$ of a specific run was sent the receiver generated a joker, let T be the set of these jokers, we have that $|T|$ and the number of tokens $< q_s, * >$ is at least $o+1$. Note that if one token in T was used by an agent then there exists an agent a_1 with a token $< q_x, l > \in Jokers$ and either: (1) there exists a token $< q_x, l >$ on agent a_2, therefore after a finite number of interactions (a_2, a_1) we have that agent a_1 will put a joker in $sending$; or (2) if token $< q_x, l >$ does not exists then it was lost during an interaction but a corresponding joker was generated.

Therefore the run for state q_s exists and thus it exists a sequence of interactions that move this run, or another, in the $sending$ buffer of q_r, therefore a_r can execute $state_{\mathcal{P}} = \delta(q_s, state_{\mathcal{P}})[1]$.

If there are only tokens $< (q_s, *), * >$ then by using the previous reasoning we can show that there is a run of such tokens, from this point on we boil down to a previous case. A similar argument shows that if there are both tokens, $< q_s, * >$ and $< (q_s, *), * >$ there is a run for at least one of them.

The only case left is if there is no agent with $state_{sim} =$ available, in this case we show that it must exists an agent a_r with $state_{\mathcal{P}} = q_r$ and a run of tokens $< q_r, * >$ or $< (q_r, *), * >$. Let us assume the contrary, when an agent switches to state pending it inserts a run of tokens $< q_r, * >$ in $sending$. This run can only disappears if is consumed by an agent, but in this case a run $< (q_r, *), * >$ is created, now if this run is consumed by another agent a_x with state q_r this implies that the run $< q_r, * >$ of a_x is still in the system.

Therefore there always exists a sequence of interactions that bring a_r to state available.

Now we have to show that each time an agent a_r sets $state_{\mathcal{P}} = \delta(q_s, state_{\mathcal{P}})[1]$, it can be paired consistently in the matching. When, a_r changes state, at time t it consumes a run $< q_s, * >$, this implies that there exists an instant $t' < t$ where an agent a_s has generated the run $< q_s, * >$ entering in state $pending$. Now a_s could exit from $pending$ only if it consumes a run $< q_s, * >, < (q_s, *), * >$ if is the run generated by a_r we have the edge of the matching, otherwise this implies that there exist another agent a_b with state q_s that generated a run $< q_s, * >$ and consumed the run of $< (q_s, *), * >$ of a_r at time t'', for the moment let us assume $b \neq r$, if a_s was pending at time t'', then, being agent anonymous, we can switch the role of a_s and a_r and match the two agents. Otherwise if a_s was not pending, we have that when a_s exited from state pending there must exists another agent a_x with state q_s in pending, therefore we can switch the role of a_s and a_x. Let us now study the case where $b = r$, in this case we have $\delta(q_s, state_{\mathcal{P}})[1] = q_s$ thus when a_r consumes the run $< (q_s, q_r), * >$ generated by itself it enters in state $\delta(q_s, state_{\mathcal{P}})[0] = q_x$, therefore we can match a_s, a_r, switching the role of a_s and a_r, since it is as they transitioned from (q_s, q_r) to (q_x, q_s). So we have shown that under GF the matching is not empty and contains infinitely many pairs.

Finally, the derived execution is GF due to Lemma 2. □

Proof of Theorem 6

In this Section we show a naming algorithm $\mathcal{N}_n$ that is used to prove Theorem 6. Theorem 6 complements the impossibility of Section III-A, showing that a minimal amount of global knowledge, the size of the system, it is enough to build a simulator. Moreover, the simulation is possible also under the stronger UO Adversary.

Naming Algorithm: $\mathcal{N}_n$*:* The following naming protocol, $\mathcal{N}_n$, uses the knowledge of n. This naming protocol is similar to the threshold protocol for IO presented in [4]. Protocol Variables: Each agent a_r has variables $my\_id = 1$,$max\_id = 1$, moreover it can access simulator $\mathcal{S}_{ID}$ by function start_sim(id) that takes as input an unique id. Simulator Protocol: When an agent a_r is the responder of an interaction and the starter has its same value for my_id it increments the value of its variable my_id. Similarly, variable max_id is updated to reflect the maximum value seen for variable my_id, a_r updates the value max_id to the maximum value between the values my_id, max_id of the initiator and its variables. When value $max\_id = n$ the agent invokes start_sim(max_id).

Lemma 3. *Let us consider a set of agents* A *running* $\mathcal{N}_{\diamond n}$ *algorithm. The system fairness is* GF *and the interaction model is* IO*. The maximum value* M *for* max_id *increases if and only if there are two agents that share the same value for* my_id*. When* $M = |A|$ *each agent has an unique stable value for* my_id*.*

Proof. If the maximum value M for max_id increased there exists an agent a that at the end of an interaction has $my\_id > max\_id = M$, but a increases my_id only when it is the responder of an interaction with an agent a' with same value for $my\_id = M$. For the other direction let us suppose that there are two agents a, a' with same value for my_id, if one of them has value $my\_id = M$ then for GF they will eventually interacts and this increase M, if they have a value $my\_id = M' < M$ they will eventually interacts and one of them, let us suppose a', will increase its value for $my\_id = M'' = M' + 1$. By definition of M there must exist another agent a'' with $my\_id = M''$, thus we can iterate this reasoning until we have two agents with equal $my\_id = M$. It remains to show that eventually all agents will assume an unique value for my_id in $[1, \ldots, |A|]$, but this is trivial by observing that given a set of agents with equal value for $my\_id = M'$ there will be interactions by them until only two agents with value $my\_id = M'$ remains. When this happen, they eventually interact and only one agent with $my\_id = M'$ remains, this agent never changes my_id, the rest derives from the initial state $my\_id = 1$ for all agents. □

The correctness of this protocol derives immediately from Lemma 3 and Theorem 5. Thus, the following Theorem holds:

Theorem 6. *Assuming* IO*, knowledge of* $|A| = n$*, and* $\Theta(\log n)$ *bits of memory, there exists a simulator for every* TW *protocol.*