A Kinetic Analysis of the Impact of Message Loss in a Residual Gossip Protocol

Simone Dallospedale[1,†], Federico Bergenti[2,*,†] and Stefania Monica[3,†]

[1]Dipartimento di Matematica e Informatica, Università degli Studi di Ferrara, Italy
[2]Dipartimento di Ingegneria dei Sistemi e delle Tecnologie Industriali, Università degli Studi di Parma, Italy
[3]Dipartimento di Scienze e Metodi dell'Ingegneria, Università degli Studi di Modena e Reggio Emilia, Italy

Abstract

Gossip protocols are widely used in large-scale multi-agent systems to achieve distributed coordination through local interactions. In practice, however, communication is often unreliable due to packet loss or intermittent connectivity, which can significantly affect protocol behavior. This paper studies the impact of message loss on the dynamics of a residual gossip protocol. Agents interact according to a well-known residual gossip mechanism, while communication unreliability is modeled by a random variable indicating the success or failure of each message. Specifically, each message is lost with a fixed probability, identical across agents and interactions, so that the sender always updates its state, whereas the receiver updates its state only if the message is successfully received. Starting from the adopted interaction rules, a kinetic description is derived to characterize the collective dynamics of the multi-agent system. Closed-form expressions for the evolution of the mean values of the state components are obtained, providing an explicit quantification of the effect of message loss on the distributed averaging process. The analysis reveals how the collective dynamics depend on the message loss probability and derives a formula to reconstruct the initial average from the asymptotic mean values. Analytical predictions are validated through simulations of large-scale multi-agent systems, showing strong agreement between the predicted and the simulated behaviors.

Keywords

Message Loss, Residual Gossip Protocol, Kinetic Theory of Multi-Agent Systems

1. Introduction and Motivation

The *distributed averaging problem* [1] is a fundamental task in the coordination of large-scale *Multi-Agent Systems (MAS)*. Given a collection of agents, the goal is to compute the average of locally held values through local interactions only, without relying on centralized coordination [1, 2]. This problem falls within the broader class of *consensus problems*, where agents iteratively exchange information to reach agreement on quantities of interest [3, 4]. In these settings, each agent is associated with a state variable, and the objective is asymptotic convergence to a common value. Distributed averaging is a specific instance of consensus in which the agreement value coincides with the average of initial states.

Consensus problems arise in diverse domains, including distributed computing and distributed artificial intelligence [5, 6, 7]. In these scenarios, centralized coordination is often impractical due to scalability limitations, communication costs, and robustness concerns. These challenges have motivated the design of decentralized algorithms to achieve global objectives through purely local interactions.

Among decentralized approaches, gossip-based algorithms provide a simple and scalable solution to the averaging problem. Their simplicity and decentralized nature make them particularly suitable for applications in *Wireless Sensor Networks (WSNs)* and the *Internet of Things (IoT)* [8, 9]. In these settings, distributed averaging naturally arises when monitoring environmental quantities such as temperature, humidity, or pollution levels [10, 9]. When centralized coordination is unavailable, gossip protocols allow agents to collaboratively compute such estimates through local information exchanges only.

WOA 2026: 27th Workshop From Objects to Agents, June 15–17, 2026, Salerno, Italy
*Corresponding author.
†These authors contributed equally.
simone.dallospedale@unife.it (S. Dallospedale); federico.bergenti@unipr.it (F. Bergenti); stefania.monica@unimore.it (S. Monica)
0000-0002-1376-8866 (S. Dallospedale); 0000-0002-4756-4765 (F. Bergenti); 0000-0001-6254-4765 (S. Monica)
© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-10-30

CEUR-WS.org/Vol-4307/paper13.pdf

A widely studied gossip protocol is the randomized pairwise gossip protocol [1], belonging to the class of symmetric gossip methods [11]. In symmetric schemes, interacting agents exchange information and update their states in a way that preserves the global average at each step. While this guarantees correctness of the distributed averaging process, bidirectional communication increases coordination complexity and communication overhead. To mitigate these limitations, gossip protocols based on one-way communication have been proposed.

Among them, the residual gossip protocol introduced in [12] incorporates an auxiliary residual variable that compensates for asymmetric exchange of information, ensuring the preservation of the global average while maintaining communication efficiency. During each interaction, the sender agent updates its state immediately after transmitting the message, whereas the receiver agent updates its state based on the information contained in the received message.

Despite their efficiency, asymmetric gossip mechanisms are inherently more sensitive to communication imperfections. In practical distributed systems, packet loss may occur due to interference, limited transmission power, or varying link quality [13, 14]. Energy constraints may further discourage the use of reliability mechanisms such as acknowledgments or retransmissions [15]. Consequently, communication failures can significantly affect both the convergence rate and the accuracy of ordinary distributed averaging protocols [16, 17].

Motivated by these considerations, this paper investigates the impact of message loss on the residual gossip protocol introduced in [12]. The analysis builds upon the framework introduced in [18]. In that work, residual gossip dynamics are studied within the *Kinetic Theory of Multi-Agent Systems (KTMAS)* assuming reliable communication. Kinetic approaches provide a powerful framework to describe the collective behavior emerging from numerous local interactions, linking local rules to the evolution of aggregate quantities, thus enabling analytical characterization of large MAS [19, 20].

In the considered setting, communication unreliability is modeled by a random variable representing message success or failure. Losses occur independently with constant probability across agents and interactions. Each interaction involves a sender and a receiver agent, and the state updates follow the residual gossip rules. In contrast to reliable communication, the sender always updates its state, whereas the receiver updates its state only upon successful message delivery. This interaction mechanism is incorporated into a Boltzmann-type kinetic equation describing the evolution of the MAS.

Within this framework, closed-form expressions for the evolution of the mean values of agents' states are derived. The analysis quantifies the effect of message loss on the collective dynamics, revealing a degradation of the distributed averaging process. Moreover, it provides a reconstruction formula that allows recovering the initial average of the MAS from the asymptotic values of the mean state components despite communication failures. The novelty of this work lies in providing a quantitative characterization of how message loss affects the collective behavior of the considered residual gossip protocol, extending prior studies that focused on convergence and correctness under ideal communication conditions.

The remainder of the paper is organized as follows. Section 2 reviews related work. Section 3 introduces the considered residual gossip protocol and formalizes the interaction rule with message losses. Section 4 presents the foundations of the adopted KTMAS framework. Section 5 reports the analytical results, while Section 6 provides experimental validation using a custom-built simulator. Finally, Section 7 concludes the paper and outlines directions for future research.

2. Related Work

Gossip protocols constitute a well-established class of decentralized protocols for distributed averaging and consensus in large-scale networks [8, 2]. Early work on symmetric gossip protocols, such as the randomized gossip algorithm [21, 22, 1], guarantees convergence to the exact average by relying on bidirectional information exchange between randomly selected pairs of agents. While this mechanism ensures correctness, the requirement for two-way communication increases coordination and messaging overhead, which may become significant, especially in large-scale or resource-constrained MAS.

Several variants of gossip protocols have been proposed to improve communication efficiency by modifying how information is disseminated across the network or by exploiting additional structural properties. For example, geographic gossip [23, 24] integrates gossiping with geographic routing, leveraging nodes' location awareness to enable long-range exchanges and mitigate slow convergence in large networks. Similarly, broadcast gossip [25] exploits the broadcast nature of wireless communications, allowing a transmitting agent to update multiple neighbors simultaneously. A comparative analysis of these and other protocols under practical constraints is provided in [26]. These studies highlight how different communication strategies and dissemination patterns can significantly affect the efficiency and convergence behavior of distributed averaging protocols.

A complementary line of research focuses instead on reducing communication costs by modifying the state update mechanism. In particular, asymmetric gossip schemes [11] are characterized by interactions in which only one of the two agents updates its state based on the information received from the other. This contrasts with symmetric gossip protocols, where both agents update their states at each interaction. While such schemes significantly reduce coordination and communication requirements, they generally do not preserve the exact average. As shown in [11], the asymptotic consensus value may deviate from the true average, although the induced random bias may vanish in the large-scale limit under suitable and mild assumptions.

The residual gossip protocol introduced in [12] addresses this limitation by introducing an auxiliary residual variable for each agent, which compensates for the imbalance created by one-way communications. This mechanism ensures preservation of the global average under reliable communication while retaining the efficiency of asymmetric schemes, making it particularly appealing for large-scale, resource-constrained MAS. Conceptually, this approach is related to the principle of mass conservation, which also underlies distributed algorithms such as the push-sum protocol proposed in [8]. In both cases, auxiliary variables compensate for the imbalance introduced by one-way exchanges; however, the specific update schemes differ, resulting in distinct mechanisms for preserving the global average.

The robustness of gossip protocols under unreliable communication has also attracted attention, although the literature explicitly addressing pairwise schemes with message loss is still relatively limited. In symmetric gossip protocols, for instance, the effect of packet loss on the preservation of the average has been analyzed in [27]. The study shows that whether the global average is preserved critically depends on the update mechanism: if both agents update their states only upon successful communication, the average is preserved; otherwise, independent updates in the presence of failed transmissions may lead to biased consensus values.

More generally, consensus under communication uncertainties has been studied in graph-based settings, where packet drops and link failures are modeled at the level of network edges and analyzed in terms of stability and convergence properties [16, 17]. While these works provide important insights into convergence guarantees, they typically do not provide an explicit characterization of the collective dynamics induced by stochastic interaction failures. Such characterization is crucial for understanding MAS-level effects of unreliable communications.

An alternative and complementary perspective is offered by kinetic modeling of MAS, where Boltzmann-type equations describe the evolution of the probability distribution of agents' states rather than individual trajectories [19]. This approach has been successfully applied to opinion dynamics and consensus processes, enabling analytical derivation of collective observables such as mean values and variances [28, 29]. More recently, kinetic approaches have been extended to analyze gossip protocols, including the residual gossip scheme, under the assumption of reliable communication [30, 18].

Building upon these results, the present work incorporates stochastic message loss directly into the residual gossip interaction mechanism and analyzes its impact within a kinetic framework. Unlike classical graph-based approaches, the adopted approach does not rely on an explicit graph structure, rather it models interactions as stochastic pairwise exchanges with no explicit relationship with the underlying communication or proximity network. This enables a direct characterization of the collective evolution of the MAS and provides analytical insights into how message loss quantitatively degrades the distributed averaging process, bridging a gap between collective interaction rules and emergent behavior of the MAS as a whole.

3. A Residual Gossip Protocol with Unreliable Communication

This section introduces the residual gossip protocol considered in this paper, by extending the original formulation in [12] and its kinetic analysis in [18] to explicitly account for stochastic message losses. The protocol operates through pairwise interactions among randomly selected agents. While the standard protocol assumes reliable communication, a more realistic scenario with unreliable channels is addressed here, where messages may be lost independently with a given probability. To capture this effect, a stochastic interaction rule is introduced: the receiver's state is updated according to a Bernoulli random variable representing the success or failure of message transmission. This approach embeds communication unreliability directly into the local dynamics and provides the basis for the kinetic analysis presented in Section 4 and Section 5.

The protocol is studied in a MAS $\mathcal{M}$ composed of $n \in \mathbb{N}_+$ agents operating under the following common and mild assumptions:

1. The number of agents is constant over time and sufficiently large to justify a statistical description of the MAS;
2. Each agent is uniquely and permanently identified by an integer in $\{1, \dots, n\}$;
3. Agents communicate through one-way messages with negligible transmission delay; and
4. Any agent can send messages to any other agent in the MAS.

At time $t \in \mathbb{R}_{\geq 0}$, each agent i is characterized by a state $(x_i(t), y_i(t)) \in \mathbb{R}^2$, where $x_i(t)$ denotes the *primary component* (or x *state component*) and $y_i(t)$ the *residual component* (or y *state component*). The protocol starts at time $t = 0$, when each agent initializes its state components. Specifically, the initial values $x_i(0)$ and $y_i(0)$ are independently drawn from uniform distributions over the intervals $[l_x, u_x)$ and $[l_y, u_y)$, respectively, where $l_x, u_x, l_y, u_y \in \mathbb{R}$, with $l_x < u_x$ and $l_y < u_y$. Let $\overline{x}_0$ and $\overline{y}_0$ denote the mean values of the primary and residual components across the agents at time $t = 0$, respectively. More generally, let $\overline{x}(t)$ and $\overline{y}(t)$ denote the corresponding means at time $t \in \mathbb{R}_{\geq 0}$, so that $\overline{x}(0) = \overline{x}_0$ and $\overline{y}(0) = \overline{y}_0$. Finally, let $\sigma_x^2(t)$, $\sigma_y^2(t)$ and $\sigma_{xy}(t)$ denote the variances and covariance of the x and y state components across the agents in $\mathcal{M}$ at time t.

The MAS evolves through pairwise interactions. Consider two distinct agents: a sender identified by $s \in \mathbb{N}_+$, with $1 \leq s \leq n$, and a receiver identified by $r \in \mathbb{N}_+$, with $1 \leq r \leq n$ and $r \neq s$. When the sender interacts with the receiver through a one-way message at time $t \in \mathbb{R}_{\geq 0}$, the standard residual gossip protocol prescribes the following state updates [18]. When acting as sender, the agent updates its state $(x_s(t), y_s(t))$ according to

$$\begin{cases} x_s'(t) = x_s(t) \\ y_s'(t) = \alpha y_s(t), \end{cases} \tag{1}$$

independently of the receiver. If the transmitted message is successfully received, the receiver updates its state $(x_r(t), y_r(t))$ according to

$$\begin{cases} x_r'(t) = x_r(t) - \gamma\big(x_r(t) - x_s(t)\big) + (1-\gamma)(1-\alpha) y_s(t) \\ y_r'(t) = y_r(t) + \gamma\big(x_r(t) - x_s(t)\big) + \gamma(1-\alpha) y_s(t). \end{cases} \tag{2}$$

The protocol parameters $\alpha \in (0,1)$ and $\gamma \in (0,1)$ denote the *residual parameter* and the *averaging parameter*, respectively. These parameters are shared among all agents in the MAS and regulate the interaction dynamics. Their role can be inferred from equations (1)-(2).

The parameter γ modulates the consensus mechanism by controlling the strength of the averaging process. Specifically, the term $\gamma(x_r(t) - x_s(t))$ drives the receiver's state $x_r(t)$ towards the sender state $x_s(t)$ by reducing their discrepancy. Neglecting the contribution of the residual $y_s(t)$, the update of the receiver's primary component simplifies to a convex combination of the pre-interaction states. This shows that larger values of γ yield a stronger attraction towards the sender's x state component.

The parameter α regulates the evolution of the residual component. When acting as sender, an agent retains a fraction α of its residual, while the remaining fraction $(1-\alpha)$ is transferred to the receiver. This

transferred quantity is then split between the receiver's primary and residual component. Specifically, a fraction $(1-\gamma)(1-\alpha)y_s(t)$ is added to the update of $x_r(t)$, while the remaining fraction $\gamma(1-\alpha)y_s(t)$ is used in the update of $y_r(t)$. Therefore, the parameters α and γ jointly determine the interaction dynamics and the redistribution of information across the two state components.

The interaction rule is inherently asymmetric, in the sense that the sender updates its state independently, whereas the receiver updates its state based on the received information. This structure enables a one-way communication model without acknowledgments or synchronization, allowing agents to initiate interactions without waiting for responses and reducing coordination overhead. The use of one-way messages also allows the assumption of instantaneous interactions, which is common in decentralized protocols, to be better satisfied. Such a design is particularly suitable for large-scale or resource-constrained MAS and aligns with the class of asymmetric gossip schemes, while preserving the exact global average under reliable communication by introducing the residual component.

In practical scenarios, communication channels may be unreliable, leading to message losses. To account for this phenomenon, a Bernoulli random variable $\chi \in \{0,1\}$ is introduced, where $\chi = 1$ denotes successful transmission and $\chi = 0$ represents message loss. The parameter $p \in [0,1]$ denotes the probability of message loss. Thus, $\chi = 1$ (successful delivery) occurs with probability $1-p$ and $\chi = 0$ (message loss) occurs with probability p. The random variable χ is assumed independent across interactions and independent of the agents' states. It is important to note that χ is not explicitly available to the agents during the execution of the protocol. Message losses are not directly observable by the agents, which always apply the deterministic update rules (1)-(2); the stochasticity introduced by χ captures communication unreliability at the modeling level.

To incorporate message losses, the effective local interaction rule modeling the dynamics of the MAS can be written as

$$\begin{cases} x_s'(t) = x_s(t) \\ y_s'(t) = \alpha y_s(t) \\ x_r'(t) = x_r(t) + \chi\big[-\gamma\big(x_r(t) - x_s(t)\big) + (1-\gamma)(1-\alpha)y_s(t)\big] \\ y_r'(t) = y_r(t) + \chi\big[\gamma\big(x_r(t) - x_s(t)\big) + \gamma(1-\alpha)y_s(t)\big], \end{cases} \tag{3}$$

where the same realization of the Bernoulli random variable χ determines the update of both receiver state components. When $\chi = 1$, the message is successfully delivered and the interaction coincides with the deterministic residual gossip update in (1)-(2). When $\chi = 0$, the message is lost and the receiver does not update its state, while the sender still performs its local update. Specifically, the sender updates its residual component while the receiver remains unchanged, so that the fraction $(1-\alpha)y_s(t)$ of information is effectively not transferred. As a result, the total quantity $\sum_{i=1}^{n}(x_i(t) + y_i(t))$ is no longer conserved at the interaction level. This lack of conservation is a key feature of the model and motivates the collective analysis developed in the following sections.

In the considered protocol, interactions occur asynchronously in time [1]. Each agent waits for a random time interval and then initiates a communication with another agent selected uniformly at random from the remaining agents. The waiting times are exponentially distributed with rate $\hat{\nu} \in \mathbb{R}_+$, so that each agent generates transmission events according to a Poisson process with rate $\hat{\nu}$. The parameter $\hat{\nu}$, referred to as the *sending rate*, represents the average number of messages sent by each agent per unit time and is shared among all agents. At the MAS level, the asynchronous transmissions induce a sequence of interaction events. Since each agent follows an independent Poisson process of rate $\hat{\nu}$, the overall interaction process can be modeled as the superposition of these n individual processes. By the superposition property of independent Poisson processes, the resulting global interaction process is still Poisson, with rate

$$\nu = n\hat{\nu}, \tag{4}$$

which represents the average number of interactions per unit time in the MAS. Consequently, the time between successive interactions is exponentially distributed with parameter ν, providing a stochastic description of the event-driven evolution of the MAS consistent with the kinetic framework briefly reviewed in Section 4.

Controller 1 Pseudocode of the agent controller for the studied residual gossip protocol

Require: $\alpha \in (0,1), \gamma \in (0,1), \hat{\nu} \in \mathbb{R}_+$
State: $x \in \mathbb{R}, y \in \mathbb{R}, \mathbf{X} \sim \mathcal{U}([l_x, u_x)), \mathbf{Y} \sim \mathcal{U}([l_y, u_y)), \mathbf{D} \sim \mathcal{E}(\hat{\nu})$

procedure UPDATE(x', y')
  $x \leftarrow x'$ ▷ Update x
  $y \leftarrow y'$ ▷ Update y
on CREATED
  $x_0 \leftarrow \mathbf{X}$ ▷ Set the initial value for x
  $y_0 \leftarrow \mathbf{Y}$ ▷ Set the initial value for y
  UPDATE(x_0, y_0) ▷ Update x and y
  SLEEP($\mathbf{D}$) ▷ Sleep for a random interval
on RECEIVED (x_s, y_s)
  $x' \leftarrow x - \gamma(x - x_s) + (1-\gamma)(1-\alpha)y_s$ ▷ Apply the interaction rule (receiver side)
  $y' \leftarrow y + \gamma(x - x_s) + \gamma(1-\alpha)y_s$
  UPDATE(x', y') ▷ Update x and y
on AWAKENED
  $r \leftarrow$ GET_RANDOM_PEER ▷ Select a random peer
  SEND($r, (x, y)$) ▷ Send a message
  $x' \leftarrow x$ ▷ Apply the interaction rule (sender side)
  $y' \leftarrow \alpha y$
  UPDATE(x', y') ▷ Update x and y
  SLEEP($\mathbf{D}$) ▷ Sleep for a random interval

Each agent can interact with any other agent, selected uniformly at random. Self-interactions are excluded by the protocol. This exclusion does not affect the global interaction rate ν, since each agent still generates transmission events according to an independent Poisson process with rate $\hat{\nu}$. The exclusion of self-interactions only modifies the receiver selection mechanism, which is performed uniformly over the remaining $n-1$ agents. Moreover, asynchronous operation ensures that simultaneous events occur with negligible probability, so that each agent is involved in at most one interaction at any given time. Consequently, an agent can process (send or receive) only a single message at a time.

Controller 1 presents the pseudocode used by each agent to implement the considered residual gossip protocol, including state initialization, random activation, peer selection, and update rules. The controller is parameterized by α, γ, $\hat{\nu}$, and the initialization intervals $[l_x, u_x)$ and $[l_y, u_y)$. It maintains the agent state composed of the primary component x and the residual component y, which are initialized by sampling $\mathbf{X}$ and $\mathbf{Y}$ from uniform distributions over the intervals $[l_x, u_x)$ and $[l_y, u_y)$, respectively. Random waiting times between interactions are generated according to an exponential random variable $\mathbf{D}$ with rate $\hat{\nu}$. At each awakening, the controller selects a peer uniformly at random from the remaining $n-1$ agents using the procedure GET_RANDOM_PEER and transmits its current state using the procedure SEND.

To explicitly model unreliable communication, message loss can be embedded in this controller by incorporating the message loss probability $p \in [0,1]$ in the SEND statement. When an agent uses SEND, the message is passed to the underlying delivery mechanism whose erratic behavior can be described using a Bernoulli random variable $\mathbf{B} \sim \mathcal{B}(1-p)$. If $\mathbf{B} = 1$, the receiver receives the message and executes the corresponding RECEIVED procedure; otherwise, the message is discarded, modeling a communication failure. This description decouples the agent logic from the stochastic network behavior: agents always apply the deterministic update rules (1)-(2), while message loss is implicitly captured by the probabilistic nature of the delivery mechanism. This description allows for a direct analytical characterization of the collective dynamics within the KTMAS framework introduced in the following section. This framework captures the essential behavior of large-scale MAS while maintaining generality and scalability.

4. The Basics of the KTMAS Framework

The analysis of large-scale MAS becomes particularly challenging when focusing on individual agent behavior. In the absence of centralized coordination, tracking the full state evolution of each agent quickly becomes impractical, especially as the system size increases. Consider, for instance, a large population of mobile robots or autonomous taxis operating in urban environments [31, 32]. In such scenarios, continuously collecting and processing detailed information about each agent, such as position, velocity, energy level, or internal state, may be computationally expensive and often infeasible.

At the same time, such a fine-grained description is often unnecessary when the objective is to understand or regulate the system at a global level. Instead, the primary focus typically lies in ensuring that the system as a whole achieves its overall objectives. For instance, a robotic swarm may be required to exhibit coordinated motion patterns, while an autonomous taxi fleet should guarantee an appropriate spatial distribution across the service area in order to effectively meet transportation demand.

Under these conditions, analyzing the collective behavior arising from agent interactions becomes more appropriate than focusing on individual trajectories [33, 34]. This shift in perspective directs the analysis toward the collective properties of the MAS. This perspective is well established in the study of complex systems, where large populations of interacting entities give rise to emergent phenomena that cannot be directly inferred from individual agent dynamics [35]. When agent states can be represented as continuous-valued variables, statistical methods offer a natural framework to describe system dynamics in terms of aggregate quantities. Although these quantities do not capture the detailed behavior of individual agents, they are sufficient to characterize the evolution and emergent properties of the MAS.

Within this broader context, kinetic modeling has emerged as a powerful mathematical paradigm for the analysis of large-scale interacting systems. In particular, the KTMAS framework [30, 36] has been proposed as a general analytical framework for the study of MAS governed by decentralized interaction protocols. From a methodological standpoint, the KTMAS framework can be viewed as a domain-specific extension of classical kinetic theory [37, 38, 39, 19], originally developed in statistical physics to describe particle systems such as rarefied gases [40, 41].

Rather than explicitly tracking the evolution of each agent, the KTMAS framework adopts a mesoscopic perspective to characterize the collective behavior of the system. In this view, the system is described through the evolution of statistical distributions over agent states. This mesoscopic representation provides a bridge between the microscopic dynamics of individual agents and the macroscopic behavior of the MAS, enabling the characterization of average behaviors and emergent patterns while preserving analytical tractability.

More generally, the KTMAS framework belongs to a wider class of kinetic approaches that have progressively extended beyond their original physical domain to encompass a variety of complex systems. These methods have been successfully applied to socio-economic processes, collective decision-making, and engineered MAS. Representative examples include the modeling of wealth distribution and economic exchanges [42, 43], the analysis of opinion dynamics [20, 44, 29, 45], and the study of traffic flow and crowd behavior [46, 39]. Such applications highlight the flexibility of kinetic frameworks in capturing emergent behaviors across heterogeneous domains. These models typically rely on Boltzmann-type integro-differential equations to describe the evolution of probability distributions driven by pairwise interactions, external inputs, and environmental factors [19].

The KTMAS framework specializes these principles to systems in which interactions are mediated by information exchange mechanisms, such as direct communication and decentralized coordination protocols [47]. This framework has been employed to analyze a variety of communication protocols for different MAS [30, 36, 18, 48], providing a formal link between local interaction rules and the resulting behavior of the MAS as a whole.

The terminology and notation adopted in Section 3 follow the conventions introduced by the KTMAS framework, thereby facilitating the formal analysis of the protocol considered in this work and its induced collective dynamics. In the KTMAS framework, the MAS is assumed to consist of a large but fixed number of agents $n \in \mathbb{N}_+$, each uniquely identified by an integer in $\{1, \dots, n\}$. The state of each agent is described by a vector $\mathbf{q}$ belonging to a set $Q \subseteq \mathbb{R}^m$, for some $m \in \mathbb{N}_+$. The set Q is

assumed to be open and contains all admissible configurations of the agent states. Agents interact autonomously through message exchanges. Each interaction involves two agents: a sender s and a receiver r, with $1 \leq s \leq n$ and $1 \leq r \leq n$. During an interaction, agent s sends a message to agent r, possibly triggering a response from r to s. Interactions occur randomly in time and are assumed to follow a Poisson point process with rate $\nu \in \mathbb{R}_+$, commonly referred to as the *interaction rate*. Each interaction is considered instantaneous and statistically independent from the others, so that no agent can participate in multiple interactions at the same time [36].

In the current formulation of the KTMAS framework, environmental forces are not explicitly considered, and therefore the evolution of the agent states is entirely determined by pairwise interactions. Each interaction is described by an interaction rule that maps the pre-interaction states $(\mathbf{q}_s, \mathbf{q}_r)$ of the two agents to their post-interaction states $(\mathbf{q}'_s, \mathbf{q}'_r)$. This rule depends on the specific protocol under study and remains unspecified within the general framework until a concrete model is introduced. The only requirement imposed by the KTMAS framework is that interaction rules must be total and locally invertible over the state space Q.

In the following, the focus is on the case $m = 2$, which corresponds to the protocol studied in this work. Accordingly, the state space is $Q = \mathbb{R}^2$. Following the standard formulation of kinetic theory, the collective state of the MAS is described by a *density function* $f : Q \times \mathbb{R}_{\geq 0} \to \mathbb{R}$. The quantity $f(x, y, t)\,\mathrm{d}x\mathrm{d}y$ represents the expected number of agents whose state at time $t \in \mathbb{R}_{\geq 0}$ lies within the infinitesimal region $(x, x + \mathrm{d}x) \times (y, y + \mathrm{d}y)$. The function f is assumed to be sufficiently regular so that the kinetic equations governing the system dynamics are well defined [36].

Under this interpretation, the function f provides a statistical description of the distribution of agent states, connecting the local states of individual agents with the collective quantities that characterize the collective dynamics of the MAS. Using this density function, several collective quantities describing the MAS can be expressed in integral form. For instance, the total number of agents in the system is

$$n = \int_Q f(x, y, t)\,\mathrm{d}x\mathrm{d}y. \tag{5}$$

The dependence on time is omitted, reflecting the assumption that the population size remains constant. Similarly, the mean value of the agents' x state components is defined as

$$\overline{x}(t) = \frac{1}{n} \int_Q x\, f(x, y, t)\,\mathrm{d}x\mathrm{d}y, \tag{6}$$

while the mean value of the agents' y state components is

$$\overline{y}(t) = \frac{1}{n} \int_Q y\, f(x, y, t)\,\mathrm{d}x\mathrm{d}y. \tag{7}$$

The introduction of the density function allows the derivation of the two fundamental equations underlying the KTMAS framework. These equations correspond to the (*strong form of the*) *Boltzmann equation* and *weak form of the Boltzmann equation* [19, 36]. The strong form describes the full time evolution of the density function, providing a detailed statistical characterization of the MAS. However, this description is often unnecessarily detailed for the analysis of large MAS. For this reason, the present work focuses on the weak form of the Boltzmann equation. This formulation allows the direct derivation of evolution equations for selected collective quantities, while remaining consistent with the underlying local dynamics. This approach is employed to obtain the analytical results presented in Section 5.

Given a sufficiently regular *test function* $\phi : Q \to \mathbb{R}$, the weak form of the Boltzmann equation can be written as [19]

$$\int_Q \frac{\partial f}{\partial t}(\mathbf{q}, t)\phi(\mathbf{q})\,\mathrm{d}^2\mathbf{q} = \frac{\beta}{2} \int_{Q^2} \big\langle \Delta(\mathbf{q}_s, \mathbf{q}_r) \big\rangle f(\mathbf{q}_s, t) f(\mathbf{q}_r, t)\,\mathrm{d}^2\mathbf{q}_r \mathrm{d}^2\mathbf{q}_s, \tag{8}$$

where $\langle \cdot \rangle$ denotes the expected value with respect to the random variables that may appear in the interaction rule. These random variables are assumed to be independent of the agent states and,

consequently, of the density function f. It is worth emphasizing that the explicit inclusion of this expectation extends the previous formulation of the KTMAS framework [30, 36], in which interaction rules were typically assumed to be deterministic. The parameter $\beta \in \mathbb{R}_+$ is defined, for large MAS, as

$$\beta = \frac{2}{n^2}\nu, \tag{9}$$

and

$$\Delta(\mathbf{q}_s, \mathbf{q}_r) = \phi(\mathbf{q}'_r) + \phi(\mathbf{q}'_s) - \phi(\mathbf{q}_r) - \phi(\mathbf{q}_s), \tag{10}$$

where $\mathbf{q}'_s$ and $\mathbf{q}'_r$ denote the post-interaction states of the sender and the receiver, respectively, as determined by the interaction rule under consideration.

This formulation is particularly useful for studying the evolution of collective quantities, as will become clearer in the following section. For instance, choosing $\phi(x, y) = x$, the left-hand side of the weak form of the Boltzmann equation (8) becomes

$$\frac{\mathrm{d}}{\mathrm{d}t}\int_Q x f(x, y, t)\,\mathrm{d}x\mathrm{d}y = n\frac{\mathrm{d}\overline{x}}{\mathrm{d}t}(t), \tag{11}$$

which transforms equation (8) into a relation governing the time derivative of the mean value of the agents' x state component. Similarly, selecting $\phi(x, y) = y$ in the left-hand side of (8) yields

$$\frac{\mathrm{d}}{\mathrm{d}t}\int_Q y f(x, y, t)\,\mathrm{d}x\mathrm{d}y = n\frac{\mathrm{d}\overline{y}}{\mathrm{d}t}(t), \tag{12}$$

thus yielding an evolution equation for the mean value of the agents' y state components.

5. Analytical Results

Consider the MAS $\mathcal{M}$ introduced in Section 3. This section derives closed-form expressions describing the time evolution of the mean values of the x and y state components of the agents in $\mathcal{M}$. The analysis is carried out within the KTMAS framework presented in Section 4, which provides a rigorous description of the collective dynamics induced by stochastic pairwise interactions. In particular, the interaction rule must be total and locally invertible over Q^2, where $Q = \mathbb{R}^2$ denotes the space of admissible agent states. The following lemma establishes that these properties hold.

Lemma 1. *For each realization of the Bernoulli variable $\chi \in \{0, 1\}$, the interaction rule (3) defines a transformation that is total and locally invertible over $\mathbb{R}^4$, for all $\alpha \in (0, 1)$ and $\gamma \in (0, 1)$.*

Proof. Agents in $\mathcal{M}$ interact according to the rule (3), which depends on the parameters $\alpha \in (0, 1)$, $\gamma \in (0, 1)$, and on the Bernoulli variable $\chi \in \{0, 1\}$. For a fixed realization of χ, the interaction rule defines a mapping

$$(x'_s, y'_s, x'_r, y'_r) = \mathbf{t}_\chi(x_s, y_s, x_r, y_r), \tag{13}$$

where (x_s, y_s) and (x_r, y_r) denote the pre-interaction states of the sender and the receiver, and (x'_s, y'_s), (x'_r, y'_r) their post-interaction states, as defined by (3). It is straightforward to verify that $\mathbf{t}_\chi$ is linear in the variables (x_s, y_s, x_r, y_r) and therefore defines a total mapping over $\mathbb{R}^4$. The Jacobian matrix of $\mathbf{t}_\chi$ is given by

$$\partial\mathbf{t}_\chi(x_s, y_s, x_r, y_r) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \alpha & 0 & 0 \\ \chi\gamma & \chi(1-\gamma)(1-\alpha) & 1-\chi\gamma & 0 \\ -\chi\gamma & \chi\gamma(1-\alpha) & \chi\gamma & 1 \end{pmatrix}, \tag{14}$$

whose determinant is

$$\det \partial\mathbf{t}_\chi(x_s, y_s, x_r, y_r) = \alpha(1 - \chi\gamma). \tag{15}$$

Since $\alpha \in (0, 1)$, $\gamma \in (0, 1)$ and $\chi \in \{0, 1\}$, it follows that $\det \partial\mathbf{t}_\chi$ is always nonzero for all admissible values of the parameters. Hence, for each realization of χ, $\mathbf{t}_\chi$ is locally invertible on $\mathbb{R}^4$. □

Before deriving the equations governing the evolution of $\overline{x}(t)$ and $\overline{y}(t)$, it is convenient to explicitly compute the expectation with respect to the Bernoulli random variable χ of the interaction term $\Delta(\mathbf{q}_s, \mathbf{q}_r)$ appearing in (8). Unlike the previous formulation of the KTMAS framework [30, 36], which assumed deterministic interaction rules, the presence of stochastic communication requires the explicit computation of this expectation. Therefore, the following lemma provides a closed-form expression for this quantity under the interaction rule (3).

Lemma 2. *When the interaction rule (3) is adopted, the expectation of the term $\Delta(\mathbf{q}_s, \mathbf{q}_r)$ in the Boltzmann equation (8) is computed with respect to the Bernoulli random variable χ as*

$$\big\langle \Delta(\mathbf{q}_s, \mathbf{q}_r) \big\rangle = \phi(\mathbf{q}_s') - \phi(\mathbf{q}_s) + (1-p)\big[\phi(\hat{\mathbf{q}}_r') - \phi(\mathbf{q}_r)\big], \tag{16}$$

where $\hat{\mathbf{q}}_r'$ denotes the receiver's post-interaction state corresponding to a successful message delivery ($\chi = 1$), which can be computed according to equation (2).

Proof. The agents in $\mathcal{M}$ interact according to (3), where the sender's post-interaction state $\mathbf{q}_s'$ is deterministic and independent of the Bernoulli random variable χ, while the receiver's post-interaction state $\mathbf{q}_r'$ depends on χ. By definition of $\Delta(\mathbf{q}_s, \mathbf{q}_r)$ in (10), its expectation with respect to χ is

$$\big\langle \Delta(\mathbf{q}_s, \mathbf{q}_r) \big\rangle = \phi(\mathbf{q}_s') + \big\langle \phi(\mathbf{q}_r') \big\rangle - \phi(\mathbf{q}_s) - \phi(\mathbf{q}_r). \tag{17}$$

Since χ is a Bernoulli random variable with $\mathbb{P}(\chi = 1) = 1 - p$ and $\mathbb{P}(\chi = 0) = p$, it follows that

$$\big\langle \phi(\mathbf{q}_r') \big\rangle = p\,\phi(\check{\mathbf{q}}_r') + (1-p)\,\phi(\hat{\mathbf{q}}_r'), \tag{18}$$

where $\check{\mathbf{q}}_r'$ and $\hat{\mathbf{q}}_r'$ denote the receiver's post-interaction states for $\chi = 0$ and $\chi = 1$, respectively. From the interaction rule (3), when $\chi = 0$ the receiver does not update its state, so that $\check{\mathbf{q}}_r' = \mathbf{q}_r$. Substituting the expression for $\langle \phi(\mathbf{q}_r') \rangle$ from equation (18) into equation (17) gives

$$\big\langle \Delta(\mathbf{q}_s, \mathbf{q}_r) \big\rangle = \phi(\mathbf{q}_s') - \phi(\mathbf{q}_s) + (1-p)\big[\phi(\hat{\mathbf{q}}_r{}') - \phi(\mathbf{q}_r)\big], \tag{19}$$

which concludes the proof. □

Building on the previous lemma, the following proposition provides a closed-form expression describing the time evolution of $\overline{y}(t)$ in $\mathcal{M}$.

Proposition 1. *On average, the mean of the y state components of the agents in $\mathcal{M}$ evolves for $t \in \mathbb{R}_{\geq 0}$ as*

$$\overline{y}(t) = \overline{y}_0 \mathrm{e}^{-\frac{\beta n}{2}(1-\alpha)(1-\gamma+\gamma p)\,t}. \tag{20}$$

Proof. The evolution of $\overline{y}(t)$ can be derived by considering the weak form of the Boltzmann equation (8) with the test function $\phi(x, y) = y$. Substituting this test function into (8) and using the definition (7), the left-hand side becomes

$$\int_Q \frac{\partial f}{\partial t}(x_s, y_s, t)\, y_s \,\mathrm{d}x_s \mathrm{d}y_s = n \frac{\mathrm{d}}{\mathrm{d}t}\overline{y}(t) \tag{21}$$

and the corresponding right-hand side is

$$\frac{\beta}{2}\int_{\mathbb{R}^4} \big\langle y_s' + y_r' - y_s - y_r \big\rangle\, f(x_s, y_s, t) f(x_r, y_r, t)\,\mathrm{d}x_s \mathrm{d}y_s \mathrm{d}x_r \mathrm{d}y_r, \tag{22}$$

where y_s' and y_r' denote the post-interaction residual components as defined by (3). Thus, the Boltzmann equation for $\phi(x, y) = y$ reads

$$\frac{\mathrm{d}}{\mathrm{d}t}\overline{y}(t) = \frac{\beta}{2n}\int_{\mathbb{R}^4} \big\langle y_s' + y_r' - y_s - y_r \big\rangle\, f(x_s, y_s, t) f(x_r, y_r, t)\,\mathrm{d}x_s \mathrm{d}y_s \mathrm{d}x_r \mathrm{d}y_r. \tag{23}$$

Using the interaction rule (3) together with equation (16), a straightforward algebraic computation yields the following result:

$$\langle y_s' + y_r' - y_s - y_r \rangle = -(1-\alpha)y_s + (1-p)\big[\gamma(x_r - x_s) + \gamma(1-\alpha)y_s\big]. \tag{24}$$

Recalling equation (5), which relates the density function to n, together with (7), which expresses a relationship between the density function and $\overline{y}(t)$, the integral in the right-hand side of (23) becomes

$$\int_{\mathbb{R}^4} \langle y_s' + y_r' - y_s - y_r \rangle \, f(x_s, y_s, t) f(x_r, y_r, t) \, \mathrm{d}x_s \mathrm{d}y_s \mathrm{d}x_r \mathrm{d}y_r = -n^2(1-\alpha)(1-\gamma+\gamma p)\, \overline{y}(t).$$

Therefore, the weak form of the Boltzmann equation for $\phi(x, y) = y$ reduces to the first-order homogeneous linear differential equation

$$\frac{\mathrm{d}}{\mathrm{d}t}\overline{y}(t) = -\frac{\beta n}{2}(1-\alpha)(1-\gamma+\gamma p)\, \overline{y}(t) \tag{25}$$

with initial condition $\overline{y}(0) = \overline{y}_0$. Solving this equation gives

$$\overline{y}(t) = \overline{y}_0 \mathrm{e}^{-\frac{\beta n}{2}(1-\alpha)(1-\gamma+\gamma p)t} \tag{26}$$

for $t \in \mathbb{R}_{\geq 0}$, which proves the proposition. □

The next proposition extends the analysis based on the KTMAS framework by providing an analytical expression for the evolution of $\overline{x}(t)$ in $\mathcal{M}$.

Proposition 2. *On average, the mean of the x state components of the agents in $\mathcal{M}$ evolves as*

$$\overline{x}(t) = \bar{x}_0 + \frac{(1-p)(1-\gamma)}{1-\gamma+\gamma p}\big[\overline{y}_0 - \overline{y}(t)\big] \tag{27}$$

for $t \in \mathbb{R}_{\geq 0}$.

Proof. The evolution of $\overline{x}(t)$ can be obtained by choosing the test function $\phi(x, y) = x$ in the weak form of the Boltzmann equation (8). Substituting this test function into (8) and using the definition (6), the left-hand side becomes

$$\int_Q \frac{\partial f}{\partial t}(x_s, y_s, t)\, x_s \, \mathrm{d}x_s \mathrm{d}y_s = n\frac{\mathrm{d}}{\mathrm{d}t}\overline{x}(t) \tag{28}$$

and the corresponding right-hand side becomes

$$\frac{\beta}{2}\int_{\mathbb{R}^4} \langle x_s' + x_r' - x_s - x_r \rangle \, f(x_s, y_s, t) f(x_r, y_r, t) \, \mathrm{d}x_s \mathrm{d}y_s \mathrm{d}x_r \mathrm{d}y_r, \tag{29}$$

where x_s' and x_r' denote the post-interaction primary components defined by (3). Therefore, the Boltzmann equation for $\phi(x, y) = x$ becomes

$$\frac{\mathrm{d}}{\mathrm{d}t}\overline{x}(t) = \frac{\beta}{2n}\int_{\mathbb{R}^4} \langle x_s' + x_r' - x_s - x_r \rangle \, f(x_s, y_s, t) f(x_r, y_r, t) \, \mathrm{d}x_s \mathrm{d}y_s \mathrm{d}x_r \mathrm{d}y_r. \tag{30}$$

Using the interaction rule (3) together with (16), a direct algebraic manipulation results in

$$\langle x_s' + x_r' - x_s - x_r \rangle = (1-p)\big[-\gamma(x_r - x_s) + (1-\gamma)(1-\alpha)y_s\big]. \tag{31}$$

Recalling equation (5), which relates the density function to n, together with (6), which expresses a relationship between the density function and $\overline{x}(t)$, the integral in the right-hand side of (30) becomes

$$\int_{\mathbb{R}^4} \langle x_s' + x_r' - x_s - x_r \rangle \, f(x_s, y_s, t) f(x_r, y_r, t) \, \mathrm{d}x_s \mathrm{d}y_s \mathrm{d}x_r \mathrm{d}y_r = n^2(1-p)(1-\alpha)(1-\gamma)\, \overline{y}(t).$$

Therefore, the Boltzmann equation for $\phi(x, y) = x$ reduces to the first-order homogeneous linear differential equation

$$\frac{\mathrm{d}}{\mathrm{d}t}\overline{x}(t) = \frac{\beta n}{2}(1-p)(1-\alpha)(1-\gamma)\,\overline{y}(t) \tag{32}$$

with initial condition $\overline{x}(0) = \overline{x}_0$. Recalling the expression for $\overline{y}(t)$ from (20), the previous differential equation can be solved to yield

$$\overline{x}(t) = \bar{x}_0 + \frac{(1-p)(1-\gamma)}{1-\gamma+\gamma p}\left[\overline{y}_0 - \overline{y}(t)\right] \tag{33}$$

for $t \in \mathbb{R}_{\geq 0}$, which proves the proposition. □

The analytical results in Proposition 1 and Proposition 2 provide a complete characterization of the collective evolution of the MAS $\mathcal{M}$. Notably, the equations (20) and (27) do not depend on the total number of agents n. Indeed, combining (4), which links the MAS-level interaction rate ν to the sending rate $\hat{\nu}$ of each agent, with the definition of β in (9), equation (20) can be rewritten as

$$\overline{y}(t) = \overline{y}_0 \mathrm{e}^{-mt}, \tag{34}$$

where

$$m = \hat{\nu}(1-\alpha)(1-\gamma+\gamma p) \tag{35}$$

is the decay rate of $\overline{y}(t)$ and is independent of n, showing that the collective dynamics are invariant with respect to the system size.

Equation (20) and equation (27) also fully characterize the asymptotic behavior of the MAS, as formalized in Corollary 1. In particular, $\overline{y}(t)$ vanishes exponentially as $t \to +\infty$, while $\overline{x}(t)$ converges to a constant value determined by the initial mean values $\overline{x}_0$ and $\overline{y}_0$, the protocol parameter γ, and the message loss probability p.

Corollary 1. *The mean of the agents' x and y state components in $\mathcal{M}$ satisfy*

$$\lim_{t\to+\infty} \overline{y}(t) = 0 \tag{36}$$

and

$$\lim_{t\to+\infty} \overline{x}(t) = \bar{x}_0 + \frac{(1-p)(1-\gamma)}{1-\gamma+\gamma p}\,\overline{y}_0. \tag{37}$$

Proof. The result follows immediately from the explicit expressions (20) and (27). In particular, $\overline{y}(t)$ decays exponentially to zero, since equation (20) can be rewritten in the equivalent form (34), where the constant m defined in (35) is strictly positive for all admissible parameter values, namely $\alpha \in (0, 1)$, $\gamma \in (0, 1)$, and $p \in [0, 1]$. The limit for $\overline{x}(t)$ is then obtained by direct substitution. □

These results highlight the dual effect of stochastic message losses. On one hand, the decay rate of $\overline{y}(t)$ is proportional to $1 - \gamma + \gamma p$, which increases with the loss probability p, leading to faster reduction of the mean of the residual components. On the other hand, information spread occurs only during successful communications, with probability $1 - p$, slowing the convergence of the agents' primary components toward the asymptotic value of $\overline{x}(t)$.

In the reliable communication scenario ($p = 0$), the total quantity $\overline{x}(t) + \overline{y}(t)$ is conserved, in agreement with previous results [18]. When $p > 0$, this conservation no longer holds. In the extreme case $p = 1$, corresponding to complete communication failure, no information spread occurs; $\overline{x}(t)$ remains constant, while $\overline{y}(t)$ decays due to sender-side updates. Although the formulas remain well-defined, this corresponds to a degenerate case in which the protocol is effectively inactive.

Finally, it is worth noting that equation (36) also shows that distributed averaging remains achievable despite message losses: if $\overline{y}_0$ is known, as is typical in practical scenarios, the initial mean $\overline{x}_0$ can be recovered from the asymptotic value of $\overline{x}(t)$. In such settings, $\overline{y}_0$ is typically fixed a priori by

appropriately initializing the agents' residual components so that their mean equals a prescribed value, whereas the objective is to estimate the unknown value of $\overline{x}_0$.

Overall, these results provide a complete characterization of the MAS $\mathcal{M}$. They show that stochastic message losses simultaneously accelerate the reduction of the mean of the residual components and therefore decrease the efficiency of information spread caused by the primary components, ultimately hindering consensus formation.

6. Experimental Results

This section presents the experiments conducted to validate and complement the analytical results derived in Section 5. In particular, the presented experimental analysis aims to assess the accuracy of the closed-form expressions governing the evolution of the mean values of the primary and residual components, as well as to investigate the impact of message losses and protocol parameters on the collective dynamics of the studied MAS.

6.1. The Simulator

A custom discrete-event simulator was developed in Python (version 3.12.13) to reproduce the stochastic interaction dynamics of the considered residual gossip protocol. The simulator is self-contained and implemented without relying on external libraries or frameworks. The simulator is publicly available at https://www.github.com/bergenti. The simulator depends on the following parameters:

1. The number of agents $n \in \mathbb{N}_+$ in the MAs;
2. The message sending rate $\hat{\nu} \in \mathbb{R}+$ for each agent;
3. The protocol parameters $\alpha \in (0, 1)$ and $\gamma \in (0, 1)$;
4. The initialization intervals $[l_x, u_x)$ and $[l_y, u_y)$;
5. The message loss probability $p \in [0, 1]$;
6. The number of independent simulation runs $K \in \mathbb{N}_+$; and
7. The time horizon $T \in \mathbb{R}_{\geq 0}$ of each simulation.

Each simulation computes the time evolution of the mean values of the x and y state components, as well as their variances and covariance, averaged over K independent runs to reduce stochastic variability. All random variables are generated using pseudo-random number generators with a fixed seed defined in the simulator code to ensure reproducibility.

A simulation starts by initializing the agents' x and y state components, which are sampled uniformly and independently from the intervals $[l_x, u_x)$ and $[l_y, u_y)$, respectively. The delay that each agent waits before sending messages is then assigned to each agent by sampling an exponential random variable with rate $\hat{\nu}$. After initialization, at each step a sender agent is chosen among the agents with minimum delay for the next interaction. A receiver agent is then chosen uniformly at random. The sender transmits a one-way message and updates its state according to (1). With probability $1 - p$, the receiver successfully receives the message and updates its state according to (2). Before the next interaction, the delay of the sender is resampled, the delays of the other agents are decreased accordingly, and the simulation time is advanced. The process continues until the maximum time T is reached and the entire procedure is then repeated for K independent runs.

6.2. Experiments

The simulator described previously is employed to conduct the experiments presented in the remainder of this section. The analysis is organized into three experiments, each corresponding to specific parameter configurations. The simulations are intended to be illustrative and do not constitute an exhaustive validation of the analytical results.

Unless otherwise specified, the following configuration is adopted. All experiments are performed with $n = 10^3$ agents. Preliminary tests with smaller systems ($n = 100$ and $n = 500$) yielded

qualitatively similar results, in agreement with the analytical observation in Section 5 that the collective dynamics is independent of the system size. The initial values of both state components are sampled uniformly from $[0, 1)$, yielding $\overline{x}_0 = 0.5$ and $\overline{y}_0 = 0.5$ as the initial mean values. The sending rate is fixed to $\hat{\nu} = 1\,\mathrm{s}^{-1}$. The number of runs is set to $K = 100$, which is sufficient to reduce stochastic variability while maintaining moderate computational cost. The simulation horizon is set to $T = 100\,\mathrm{s}$, which allows the asymptotic regime to be observed. Each simulated MAS is denoted by $\mathcal{M}_{\alpha,\gamma,p}$.

Overall, the experiments presented in this section confirm the validity of the analytical model and highlight the distinct roles of communication unreliability and protocol parameters in shaping both the transient dynamics and the asymptotic behavior of a MAS in which agents exchange messages according to the considered residual gossip protocol.

Experiment 1: Validation of analytical expressions. The accuracy of the analytical expressions derived in Propositions 1 and 2 is first assessed. To this end, the MAS $\mathcal{M}_{0.3,0.5,p}$ is considered under both reliable ($p = 0$) and unreliable ($p = 0.2$) communication. Figures 1a and 1b compare the simulated trajectories with the corresponding analytical solutions (20) and (27). In both cases, the empirical results closely match the theoretical predictions over the entire time horizon, showing that the kinetic model accurately captures the collective evolution of the system even in the presence of stochastic message losses. The comparison between the two scenarios highlights the effect of communication unreliability on the asymptotic state of $\overline{x}(t)$. In the lossless case, the mean converges to $\overline{x}_0 + \overline{y}_0 = 1$, while for $p = 0.2$ it converges to a smaller value, consistent with (36). This shows that message losses shift the consensus value away from the lossless case.

An additional experiment is performed to assess the robustness of the analytical expressions under different initial conditions. The same system $\mathcal{M}_{0.3,0.5,p}$ is considered, but with initial x and y state components drawn uniformly from the intervals $[1, 3)$ and $[-2, 0)$, respectively. Thus, the associated initial mean values are $\overline{x}_0 = 2$ and $\overline{y}_0 = -1$. Figures 2a and 2b show the time evolution of $\overline{x}(t)$ and $\overline{y}(t)$ for both reliable ($p = 0$) and unreliable ($p = 0.2$) communication.

The results show that the analytical expressions (20) and (27) accurately predict the system dynamics even when the initial means differ from the default configuration. Specifically, the mean of the residual components $\overline{y}(t)$ still decays exponentially regardless of the initial distribution, while the mean of the primary components $\overline{x}(t)$ converges to the asymptotic value predicted by (36). This confirms that the kinetic model captures the collective behavior of the MAS independently of the initial state distribution, highlighting the generality and robustness of the analytical framework.

Experiment 2: Effect of message loss probability. The role of the message loss probability p on both transient and asymptotic dynamics is then investigated. Simulations are performed for $\mathcal{M}_{0.4,0.9,p}$ with $p \in \{0, 0.2, 0.4, 0.6, 0.8\}$. As shown in Figure 3a, increasing p produces two main effects. First, the asymptotic value of $\overline{x}(t)$ progressively shifts toward the initial mean $\overline{x}_0$, in agreement with the analytical characterization. Second, for larger values of p, the mean $\overline{x}(t)$ reaches its asymptotic value more quickly in terms of time.

However, this apparent acceleration does not correspond to faster consensus formation. On the contrary, the variance of the x state components $\sigma_x^2(t)$ reported in Figure 3b indicates that agreement among agents is reached more slowly as p increases. This reflects the reduced effectiveness of information spread: although the collective mean evolves rapidly, the local alignment process is hindered by frequent communication failures.

Experiment 3: Effect of the model parameters. The role of the protocol parameters α and γ is finally analyzed by considering the MAS $\mathcal{M}_{\alpha,\gamma,0.5}$ in two complementary cases. First, α is fixed to 0.6 and γ is varied in $\{0.1, 0.3, 0.5, 0.7, 0.9\}$. As shown in Figure 4a, the parameter γ strongly affects the asymptotic value of $\overline{x}(t)$, confirming its role in determining the steady-state behavior predicted by (36). Larger values of γ lead to asymptotic values closer to the initial mean $\overline{x}_0$. Then, γ is fixed to 0.1 and α is varied in $\{0.1, 0.3, 0.5, 0.7, 0.9\}$.

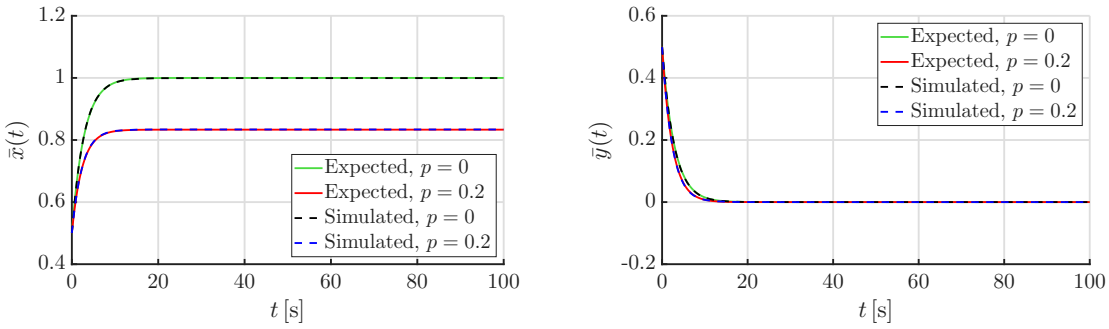

(a) (b)

Figure 1: Experiment 1 (First initial condition) – Time evolution of the mean values of the x (a) and y (b) state components for the MAS $\mathcal{M}_{0.3,0.5,p}$. Simulation results (dashed lines) are compared with analytical predictions (solid lines) for $p = 0$ and $p = 0.2$.

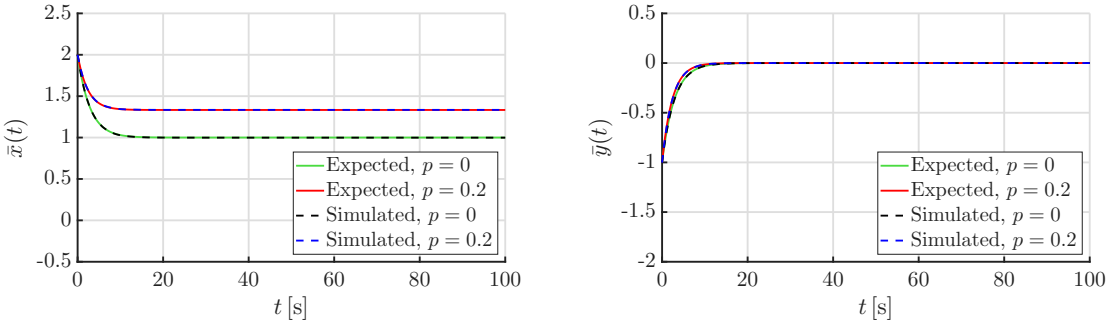

(a) (b)

Figure 2: Experiment 1 (Second initial condition) – Time evolution of the mean values of the x (a) and y (b) state components for the MAS $\mathcal{M}_{0.3,0.5,p}$ with $\overline{x}_0 = 2$ and $\overline{y}_0 = -1$. Simulation results (dashed lines) are compared with analytical predictions (solid lines) for $p = 0$ and $p = 0.2$.

The results in Figure 4b show that α does not affect the asymptotic value of $\overline{x}(t)$, but significantly influences the convergence speed. In particular, larger values of α lead to slower dynamics, in agreement with the decay rate derived in equation (35).

7. Conclusion and Future Work

This paper documents a study of the impact of unreliable communication in a residual gossip protocol. By modeling pairwise interactions through the KTMAS framework, a description of the collective dynamics is derived, explicitly accounting for stochastic message losses. Closed-form expressions describing the time evolution of the mean values of the agents' state components are obtained. The analysis shows that the mean of the residual components decays exponentially to zero for all admissible parameter values, while the mean of the primary component is driven by this decay mechanism and converges to a constant value. Moreover, message losses affect not only the convergence rate but also the asymptotic value of the mean of the primary components, leading to a deviation from the value achieved under reliable communication.

The experimental results consistently support the analytical findings. The close agreement between simulated and theoretical trajectories confirms that the proposed model accurately captures the collective behavior induced by the protocol. Furthermore, the results highlight a nontrivial effect of communication unreliability: increasing the message loss probability slows down convergence at the agent level, despite an apparently faster evolution of the mean values.

Overall, the obtained analytic results provide an accurate description of residual gossip dynamics, capturing the emergent behavior of the system under realistic communication conditions. In this sense,

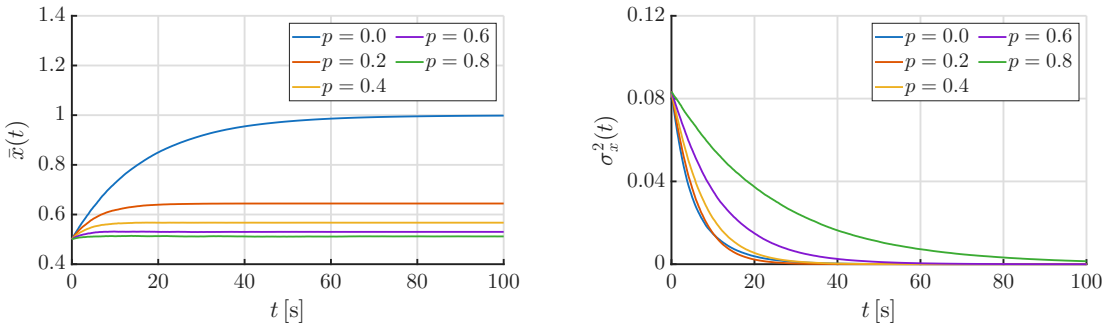

Figure 3: Experiment 2 – Time evolution of the mean (a) and variance (b) of the x state components for the MAS $\mathcal{M}_{0.4,0.9,p}$ under different message loss probabilities $p \in \{0, 0.2, 0.4, 0.6, 0.8\}$.

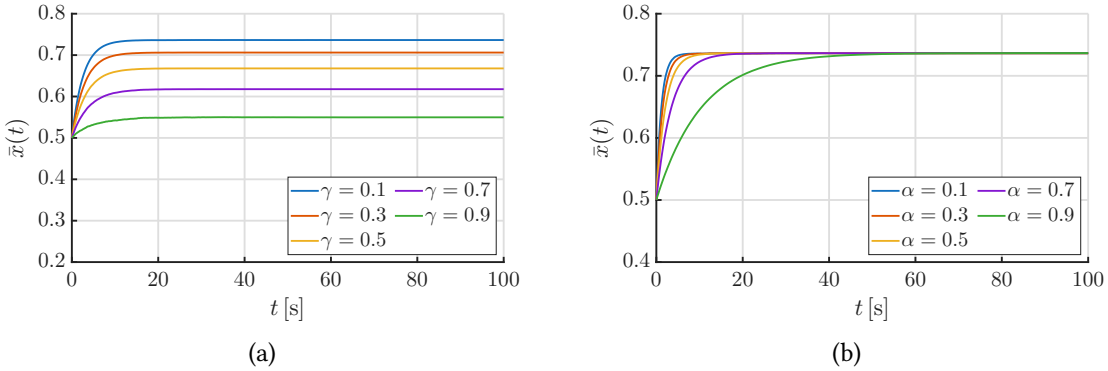

Figure 4: Experiment 3 – Time evolution of the mean of the x state components in the MAS $\mathcal{M}_{\alpha,\gamma,0.5}$ for several distinct parameter values: (a) $\gamma \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$ with fixed $\alpha = 0.6$, and (b) $\alpha \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$ with fixed $\gamma = 0.1$.

these results complement existing studies on gossip protocols, which typically focus on convergence guarantees or performance metrics, by offering a detailed characterization of the collective dynamics under unreliable communication.

Several directions for future research can be identified. A natural extension is the derivation of analytical equations for the variances and covariances of the agents' state components, in order to obtain a more complete statistical description of the system. Another promising direction is the investigation of heterogeneous communication settings, where the message loss probability depends on individual agents or groups of agents, thus enabling the incorporation of network structure into the kinetic model. Beyond these theoretical extensions, the proposed framework could also support the analysis of specific MAS, for example by inferring the message loss probability from their observed collective dynamics. These directions would further enhance the applicability of the proposed framework and deepen the understanding of gossip-based protocols operating under unreliable communication.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT in order to: Grammar and spelling check, Paraphrase and reword, Improve writing style. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] S. Boyd, A. Ghosh, B. Prabhakar, D. Shah, Randomized gossip algorithms, IEEE Transactions on Information Theory 52 (2006) 2508–2530. doi:10.1109/TIT.2006.874516.

[2] F. Fagnani, S. Zampieri, Randomized consensus algorithms over large scale networks, IEEE Journal on Selected Areas in Communications 26 (2008) 634–649. doi:10.1109/JSAC.2008.080506.

[3] M. Wooldridge, An Introduction to MultiAgent Systems, 2nd ed., Wiley Publishing, 2009.

[4] W. Ren, R. Beard, E. Atkins, A survey of consensus problems in multi-agent coordination, in: Proceedings of the 2005, American Control Conference, 2005, volume 3, IEEE, 2005, pp. 1859–1864. doi:10.1109/ACC.2005.1470239.

[5] Y. Xu, E. Liao, P. Scerri, B. Yu, M. Lewis, K. Sycara, Towards flexible coordination of large scale multi-agent teams, in: P. Scerri, R. Vincent, R. Mailler (Eds.), Coordination of Large-Scale Multiagent Systems, Springer, Boston, MA, 2006, pp. 287–309. doi:10.1007/0-387-27972-5_14.

[6] M. Jelasity, A. Montresor, O. Babaoglu, Gossip-based aggregation in large dynamic networks, ACM Transactions on Computer Systems 23 (2005) 219–252. doi:10.1145/1082469.1082470.

[7] A. F. Garcia, C. J. P. de Lucena, F. Zambonelli, A. Omicini, J. Castro (Eds.), Software Engineering for Large-Scale Multi-Agent Systems, Research Issues and Practical Applications, volume 2603 of *Lecture Notes in Computer Science*, Springer, Berlin, 2003. doi:10.1007/3-540-35828-5.

[8] D. Kempe, A. Dobra, J. Gehrke, Gossip-based computation of aggregate information, in: 44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings., SFCS-03, IEEE, 2003, pp. 482–491. doi:10.1109/SFCS.2003.1238221.

[9] A. G. Dimakis, S. Kar, J. M. F. Moura, M. G. Rabbat, A. Scaglione, Gossip algorithms for distributed signal processing, Proceedings of the IEEE 98 (2010) 1847–1864. doi:10.1109/JPROC.2010.2052531.

[10] R. Olfati-Saber, J. A. Fax, R. M. Murray, Consensus and cooperation in networked multi-agent systems, Proceedings of the IEEE 95 (2007) 215–233. doi:10.1109/JPROC.2006.887293.

[11] F. Fagnani, S. Zampieri, Asymmetric randomized gossip algorithms for consensus, IFAC Proceedings Volumes 41 (2008) 9051–9056. doi:10.3182/20080706-5-KR-1001.01528, 17th IFAC World Congress.

[12] C. Asensio-Marco, B. Beferull-Lozano, Fast average gossiping under asymmetric links in WSNS, in: Proceedings of the 22nd European Signal Processing Conference (EUSIPCO 2014), IEEE, 2014, pp. 131–135.

[13] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, E. Cayirci, Wireless sensor networks: a survey, Computer Networks 38 (2002) 393–422. doi:10.1016/S1389-1286(01)00302-4.

[14] A. Dâmaso, N. Rosa, P. Maciel, Reliability of wireless sensor networks, Sensors 14 (2014) 15760–15785. doi:10.3390/s140915760.

[15] G. Anastasi, M. Conti, M. Di Francesco, A. Passarella, Energy conservation in wireless sensor networks: A survey, Ad Hoc Networks 7 (2009) 537–568. doi:10.1016/j.adhoc.2008.06.003.

[16] F. Fagnani, S. Zampieri, Average consensus with packet drop communication, SIAM Journal on Control and Optimization 48 (2009) 102–133. doi:10.1137/060676866.

[17] S. Kar, J. M. F. Moura, Distributed consensus algorithms in sensor networks with imperfect communication: Link failures and channel noise, IEEE Transactions on Signal Processing 57 (2009) 355–369. doi:10.1109/TSP.2008.2007111.

[18] F. Bergenti, S. Monica, G. Tamboroni, F. Zambonelli, A study of the dynamics of the average in a residual gossip protocol, in: R. Collier, A. Ricci, V. Nallur, S. Burattini, A. Omicini (Eds.), Proceedings of the 21st European Conference on Multi-Agent Systems (EUMAS 2024), volume 15685 of *Lecture Notes in Computer Science*, Springer, Cham, 2025, pp. 255–273. doi:10.1007/978-3-031-93930-3_15.

[19] L. Pareschi, G. Toscani, Interacting Multiagent Systems: Kinetic Equations and Montecarlo Methods, Oxford University Press, Oxford, 2013.

[20] G. Toscani, Kinetic models of opinion formation, Communications in Mathematical Sciences 4 (2006) 481–496. doi:10.4310/CMS.2006.v4.n3.a1.

[21] S. Boyd, A. Ghosh, B. Prabhakar, D. Shah, Analysis and optimization of randomized gossip algorithms, in: 2004 43rd IEEE Conference on Decision and Control (CDC) (IEEE Cat. No.04CH37601), volume 5, IEEE, 2004, pp. 5310–5315. doi:10.1109/CDC.2004.1429652.

[22] S. Boyd, A. Ghosh, B. Prabhakar, D. Shah, Gossip algorithms: design, analysis and applications, in: Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies., volume 3, IEEE, 2005, pp. 1653–1664. doi:10.1109/INFCOM.2005.1498447.

[23] A. Dimakis, A. Sarwate, M. Wainwright, Geographic gossip: efficient aggregation for sensor networks, in: 2006 5th International Conference on Information Processing in Sensor Networks, IPSN '06, ACM Press, 2006, pp. 69–76. doi:10.1145/1127777.1127791.

[24] A. D. G. Dimakis, A. D. Sarwate, M. J. Wainwright, Geographic gossip: Efficient averaging for sensor networks, IEEE Transactions on Signal Processing 56 (2008) 1205–1216. doi:10.1109/TSP.2007.908946.

[25] T. C. Aysal, M. E. Yildiz, A. Scaglione, Broadcast gossip algorithms, in: 2008 IEEE Information Theory Workshop, IEEE, 2008, pp. 343–347. doi:10.1109/ITW.2008.4578682.

[26] M. Kenyeres, J. Kenyeres, Comparative study of distributed consensus gossip algorithms for network size estimation in multi-agent systems, Future Internet 13 (2021) 134. doi:10.3390/fi13050134.

[27] G. Shi, M. Johansson, K. H. Johansson, Randomized gossiping with unreliable communication: Dependent or independent node updates, in: 2012 IEEE 51st IEEE Conference on Decision and Control (CDC), IEEE, 2012, pp. 4846–4851. doi:10.1109/CDC.2012.6426729.

[28] S. Monica, F. Bergenti, An analytic study of opinion dynamics in multi-agent systems with additive random noise, in: G. Adorni, S. Cagnoni, M. Gori, M. Maratea (Eds.), AI*IA 2016 Advances in Artificial Intelligence, Springer International Publishing, Cham, 2016, pp. 105–117. doi:10.1007/978-3-319-49130-1_9.

[29] S. Monica, F. Bergenti, An analytic study of opinion dynamics in multi-agent systems, Computers & Mathematics with Applications 73 (2017) 2272–2284. doi:10.1016/j.camwa.2017.03.008.

[30] S. Monica, F. Bergenti, F. Zambonelli, Toward a kinetic framework to model the collective dynamics of multi-agent systems, in: Proceedings of the 11th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA 2022), volume 13703 of *Lecture Notes in Computer Science*, Springer, Berlin, 2022, pp. 165–180. doi:10.1007/978-3-031-19759-8_11.

[31] S. Mariani, G. Cabri, F. Zambonelli, Coordination of autonomous vehicles: Taxonomy and survey, ACM Computing Surveys 54 (2021) 1–33. doi:10.1145/3431231.

[32] M. Brambilla, E. Ferrante, M. Birattari, M. Dorigo, Swarm robotics: A review from the swarm engineering perspective, Swarm Intelligence 7 (2013) 1–41. doi:10.1007/s11721-012-0075-2.

[33] J. Hillston, J. Pitt, M. Wirsing, F. Zambonelli, Collective Adaptive Systems: Qualitative and Quantitative Modelling and Analysis (Dagstuhl Seminar 14512), Dagstuhl Reports 4 (2015) 68–113. doi:10.4230/DagRep.4.12.68.

[34] D. H. Wolpert, K. Tumer, Collective intelligence, data routing and Braess' paradox, Journal of Artificial Intelligence Research 16 (2002) 359–387. doi:10.1613/JAIR.995.

[35] Z. Li, C. H. Sim, M. Y. Hean Low, A survey of emergent behavior and its impacts in agent-based systems, in: 2006 4th IEEE International Conference on Industrial Informatics, IEEE, 2006, pp. 1295–1300. doi:10.1109/INDIN.2006.275846.

[36] S. Monica, F. Bergenti, F. Zambonelli, A kinetic approach to investigate the collective dynamics of multi-agent systems, International Journal on Software Tools for Technology Transfer 25 (2023) 693–705. doi:10.1007/s10009-023-00724-z.

[37] N. Bellomo, D. Burini, G. Dosi, L. Gibelli, D. Knopoff, N. Outada, P. Terna, M. E. Virgillito, What is life? A perspective of the mathematical kinetic theory of active particles, Mathematical Models and Methods in Applied Sciences 31 (2021) 1821–1866. doi:10.1142/S0218202521500408.

[38] A. Bellouquid, M. Delitala, Mathematical Modeling of Complex Biological Systems, Modeling and Simulation in Science, Engineering and Technology, Springer, Berlin, 2006.

[39] G. Albi, N. Bellomo, L. Fermo, S.-Y. Ha, J. Kim, L. Pareschi, D. Poyato, J. Soler, Vehicular traffic,

crowds, and swarms: From kinetic theory and multiscale methods to applications and research perspectives, Mathematical Models and Methods in Applied Sciences 29 (2019) 1901–2005. doi:10.1142/S0218202519500374.

[40] C. Cercignani, The Boltzmann Equation and Its Applications, volume 67 of *Applied Mathematical Sciences*, Springer, New York, 1988. doi:10.1007/978-1-4612-1039-9.

[41] C. Bianca, C. Dogbe, On the boltzmann gas mixture equation: Linking the kinetic and fluid regimes, Communications in Nonlinear Science and Numerical Simulation 29 (2015) 240–256. doi:10.1016/j.cnsns.2015.05.015.

[42] S. Cordier, L. Pareschi, G. Toscani, On a kinetic model for a simple market economy, Journal of Statistical Physics 120 (2005) 253–277. doi:10.1007/s10955-005-5456-0.

[43] L. Pareschi, G. Toscani, Wealth distribution and collective knowledge: A Boltzmann approach, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 372 (2014) 20130396. doi:10.1098/rsta.2013.0396.

[44] G. Albi, E. Calzola, G. Dimarco, A data-driven kinetic model for opinion dynamics with social network contacts, European Journal of Applied Mathematics 36 (2025) 264–290. doi:10.1017/S0956792524000068.

[45] S. Monica, F. Bergenti, Opinion dynamics in multi-agent systems: Selected analytic models and verifying simulations, Computational & Mathematical Organization Theory 23 (2017) 423–450. doi:10.1007/s10588-016-9235-z.

[46] M. Günther, A. Klar, T. Materne, R. Wegener, An explicitly solvable kinetic model for vehicular traffic and associated macroscopic equations, Mathematical and Computer Modelling 35 (2002) 591–606. doi:10.1016/S0895-7177(02)80023-1.

[47] F. Bergenti, A. Ricci, Three approaches to the coordination of multiagent systems, in: Proceedings of the 2002 ACM Symposium on Applied Computing, SAC '02, Association for Computing Machinery, New York, 2002, pp. 367–372. doi:10.1145/508791.508861.

[48] F. Bergenti, S. Monica, Experimental evaluation of a decentralized protocol for emergent heartbeat synchronization, in: Proceedings of the 2025 7th International Conference on Distributed Artificial Intelligence, DAI '25, Association for Computing Machinery, New York, 2025, pp. 19–27. doi:10.1145/3772429.3772432.