

A Poisson Jump-driven SDE Approach to Distributed Gradient Descent with Sparse Communication

Marc Weber, John Paul Strachan and Christian Ebenbauer

Abstract—To bridge the gap between idealised communication models and the stochastic reality of networked systems, we introduce a framework for embedding asynchronous communication directly into algorithm dynamics using stochastic differential equations (SDE) driven by Poisson Jumps. We apply this communication-aware design to the continuous-time gradient flow, yielding a distributed algorithm where updates occur via sparse Poisson events. Our analysis establishes communication rate bounds for asymptotic stability and, crucially, a higher, yet sparse, rate that provably any desired exponential convergence performance slower than the nominal, centralized flow. These theoretical results, shown for unconstrained quadratic optimisation, are validated by a numerical simulation.

I. INTRODUCTION

Distributed optimization has emerged as a critical necessity in numerous modern applications, spanning large-scale machine learning paradigms like Federated Learning [1], resource allocation in smart grids [2], and cooperative control in multi-agent systems and robotics [3]. The fundamental goal is to minimize a global objective function based on data or cost functions dispersed across multiple computational units or agents. Inherently, the resolution of this problem requires collaboration, necessitating communication among the distributed participants to share local information, such as gradients, parameter updates, or consensus states.

A major discrepancy often arises, however, between the theoretical models underlying many distributed optimization algorithms and the practical realities of their deployment. While the literature on distributed optimization under communication constraints is vast and diverse, much of the influential literature, particularly on foundational algorithms like Distributed Gradient Descent (DGD) and ADMM, relies on assumptions of idealized communication links [4], [5]. These assumptions frequently include perfect synchronization, fixed-clock communication rates, negligible latency, and directed or even continuous (or infinitely fast) data transfer [1], [4], [5]. Although recent advancements, notably in event-triggered optimization, address stochasticity and asynchronicity, these approaches often design the update rule first and then impose communication scarcity based on the state of the optimization (e.g., triggering a message when the local gradient significantly changes) [6], [7], [8], [9]. To our knowledge, strategies primarily focus on adapting algorithms to scarcity, rather than designing flows where channel limitations are the native drivers of the dynamics, especially in continuous-time settings. These requirements remain costly in terms of energy consumption and are prone to performance degradation under real-world

network congestion and do not consider or provide a-priori limitations on the communication channels.

This need for communication-aware design is amplified when considering modern computing architectures. Low-level distributed systems, such as processor cores connected by high-speed, asynchronous interconnects (like AMD’s Infinity Fabric [10]), and specialized systems like neuromorphic computing architectures, inherently feature a sparse, stochastic, and event-driven communication model [11], [12], [13]. In these systems, information exchange is a costly spike or packet rather than a continuous stream, and even state-of-the-art interconnects operate on time-scales much slower than efficient processing tiles [14]. Therefore, for an algorithm to be truly efficient and applicable across diverse distributed domains—from hardware fabrics to wide-area networks—it must be communication-aware from its inception, where the channel constraints dictate the structure of the mathematical flow.

To address this challenge, we propose a framework employed in control systems and networked stability analysis [15], [16], focusing on continuous-time dynamics driven by random jumps. Specifically, the random, sparse, and intermittent nature of physical communication channels and event-driven computing architectures—such as packet drops, irregular clock skews, or hardware-enforced intermittent updates—is naturally captured within this paradigm of SDEs driven by Poisson Jumps [17], [15], a framework related to piecewise deterministic Markov processes and hybrid systems modelling [18], [19]. As pioneered, for example, in the work of Brockett [17] and further developed in areas like channel selection and networked control, this framework treats the communication channel not as a fixed parameter, but as a dynamic, stochastic process integrated directly into the continuous flow that dictates the timing and magnitude of information flow [20]. This powerful conceptual tool allows us to define mathematically precise conditions for stability and performance despite high channel uncertainty, forming the theoretical core of our algorithm design.

In this paper, we leverage Poisson-driven SDE modelling to rethink the design of distributed optimization algorithms, moving beyond simply analyzing the effect of noise after the algorithm is designed. Our objective is to design algorithms where the communication strategy is an integrated, efficient component of the flow. Our primary contributions are fourfold:

Framework for Asynchronous, Sparse Design: We introduce the Poisson Jump-driven SDE modelling framework in the area of distributed optimization that uses continuous-time dynamics for the individual computation units and incorporates inter-unit communication only through stochastic Poisson-Jump events. This framework allows for the formal analysis and design of distributed algorithms under native asynchronous, sparse, and even unstable communication dynamics.

Distributed Poisson-Jump Gradient Flow: We illustrate this concept by taking the classic gradient descent flow [4] and

M. Weber and C. Ebenbauer are with the Chair of Intelligent Control Systems, RWTH Aachen University, Aachen, Germany weber@ebenbauer@ic.rwth-aachen.de, J. P. Strachan is with the Peter-Gruenberg-Institute (PGI-14), Forschungszentrum Juelich GmbH, Juelich, Germany j.strachan@fz-juelich.de.

formulating its distributed counterpart, where local gradient information is aggregated between agents via sparse, asynchronous Poisson-Jump driven communication. The resulting system is modeled as a set of coupled Poisson-driven SDEs.

Communication Rate for Stability: We establish a sufficient communication rate (related to the intensity of the Poisson process) under which the distributed optimization algorithm retains asymptotic stability (in the means square sense) toward the optimal solution. This result provides a rigorous, quantifiable bound on the intermittent, random connectivity required for stability.

Communication Efficiency and desired Rates: Crucially, we identify higher, yet still sparse range of communication rates, under which the distributed algorithm achieves any exponential convergence rate up to the one of the nominal, fully connected Gradient Descent flow. While this results gives a formal characterization of the efficiency of the communication scheme, our experiments demonstrate that exceeding the communication rate for stability does not improve the worst-case convergence rate, thus defining a target for maximally efficient, minimal-communication algorithm design.

II. PRELIMINARIES

To formalise the sparse and asynchronous communication model central to this work, we first establish the necessary mathematical framework. We begin with a brief definition of the graph-theoretic notation used to describe the network topology, followed by the stochastic process model that governs the system dynamics.

A. Graph-Theoretic Notation

We model a network of communicating agents as a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. The set of nodes $\mathcal{V} = \{\nu_1, \dots, \nu_n\}$ represents the n distributed computational units or agents. The set of edges $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ defines the communication links, where an ordered pair $(j, i) \in \mathcal{E}$ denotes a directed channel from node ν_j to node ν_i . Consequently, ν_j is an in-neighbour of ν_i , and the set of all such in-neighbours is $\mathcal{N}_i = \{\nu_j \in \mathcal{V} \mid (j, i) \in \mathcal{E}\}$

B. Poisson Jump-driven SDEs

We model the state evolution of agents and channels using SDEs driven by Poisson jumps [15], [18]. This approach naturally captures the sporadic, event-driven nature of communication.

Consider a complete probability space $(\Omega, \mathcal{F}, P)$, equipped with a filtration $\{\mathcal{F}_t\}_{t \geq 0}$ satisfying the usual conditions. We model the system state as a $\{\mathcal{F}_t\}$ -adapted stochastic process $\mathbf{x}(t, \omega) : T \times \Omega \rightarrow \mathbb{R}^n$, taking values in $\mathbb{R}^n$. Consistent with the definition of the solutions in the Itô sense, we write its dynamics in differential form as

$$d\mathbf{x}(t, \omega) = f(\mathbf{x}(t, \omega)) dt + g(\mathbf{x}(t, \omega)) dN(t, \omega), \quad (1)$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is the continuous drift vectorfield and $g : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is the jump map. The term $N(t, \omega)$ represents a standard $\{\mathcal{F}_t\}$ -Poisson counting process with rate λ . This model (1) yields a sample path $\mathbf{x}(t, \omega_0)$ for a fixed outcome $\omega_0 \in \Omega$, that is càdlàg, i.e. right-continuous with left limits, which is denoted as

$$\mathbf{x}(t, \omega_0) = \mathbf{x}(t_0, \omega_0) + \int_{t_0}^t f(\mathbf{x}(\tau, \omega_0)) d\tau + \int_{t_0}^t g(\mathbf{x}(\tau, \omega_0)) dN(\tau, \omega_0).$$

The interpretation of these pathwise dynamics, as illustrated in Figure 1, consists of two key behaviours (see [15], [18] for details):

Continuous Flow Between jump (event) times of the Poisson process $N(t, \omega)$, the system evolves deterministically according to the ordinary differential equation $\dot{\mathbf{x}}(t) = f(\mathbf{x}(t))$.

Discrete Jumps At each jump time $t_i(\omega_0)$, the state experiences an instantaneous, discrete shift $\Delta \mathbf{x}(t_i, \omega_0) := \mathbf{x}(t_i, \omega_0) - \mathbf{x}(t_i^-, \omega_0) := g(\mathbf{x}(t_i^-, \omega_0))$ where $\mathbf{x}(t^-, \omega) = \lim_{s \nearrow t} \mathbf{x}(s, \omega)$ denotes the left-limit of the state, representing the value immediately preceding the jump.




Fig. 1: Sample path of the Poisson jump-driven SDE, showing the continuous evolution of $\mathbf{x}(t, \omega_0)$ interrupted by discrete jumps.

For a function $V : \mathbb{R}^n \rightarrow \mathbb{R}$, its evolution under (1) is governed by the *Itô Differentiation Rule*

$$dV(\mathbf{x}(t, \omega)) = \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}(t, \omega)) f(\mathbf{x}(t, \omega)) dt + [V(\mathbf{x}(t, \omega) + g(\mathbf{x}(t, \omega))) - V(\mathbf{x}(t, \omega))] dN(t, \omega). \quad (2)$$

For the stability analysis in Section III, we move from the pathwise Itô rule to the expected evolution of a Lyapunov function candidate $V(\mathbf{x})$. For a sufficiently smooth function $V : \mathbb{R}^n \rightarrow \mathbb{R}$, this is given by $\frac{d}{dt} \mathbb{E}[V(\mathbf{x}(t, \omega))] = \mathbb{E}[\mathcal{L}V(\mathbf{x}(t, \omega))]$ with

$$\mathcal{L}V(\mathbf{x}) = \frac{\partial V}{\partial \mathbf{x}}(\mathbf{x}) f(\mathbf{x}) + \lambda[V(\mathbf{x} + g(\mathbf{x})) - V(\mathbf{x})].$$

This operator describes the expected rate of change of $V(\mathbf{x})$ of a sample path of (1) passing through $\mathbf{x}$. This forms the basis for the Lyapunov stability analysis via Lemma 1.

C. Stability Concepts

The following definition and theorem are adapted from [21, Definition 4.1, Theorem 4.4] with an extension for practical stability taken from the concept of ISS-stability [22, Theorem 3.4] for constant inputs to characterize the stability properties obtained in this work.

Definition 1 (Global Practical Uniform Exponential Stability in mean-square sense) A point $\mathbf{x}^* \in \mathbb{R}^n$ is said to be *globally practically uniformly exponentially stable in the mean square sense* for a SDE system (1), if there exist constants $\alpha, \beta \in (0, \infty)$ and $\gamma \in [0, \infty)$ such that for any initial condition $\mathbf{x}(t_0) = \mathbf{x}^0 \in \mathbb{R}^n$, the inequality

$$\mathbb{E}[\|\mathbf{x}(t) - \mathbf{x}^*\|^2] \leq \alpha \|\mathbf{x}^0 - \mathbf{x}^*\|^2 e^{-\beta(t-t_0)} + \gamma \quad (3)$$

holds for all $t \in [t_0, \infty)$.

If $\gamma = 0$, then the attribute *practically* can be dropped. We characterize this stability concept by means of a Lyapunov-like theorem.

Lemma 1 (Global Practical Uniform Exponential Stability in mean-square sense) Consider the SDE with jumps (1). Assume there exists a function $V : \mathbb{R}^n \rightarrow \mathbb{R}$ and positive constants c_1, c_2, c_3, γ' such that

- (1) $c_1 \|\mathbf{x} - \mathbf{x}^*\|^2 \leq V(\mathbf{x}) \leq c_2 \|\mathbf{x} - \mathbf{x}^*\|^2$ for all $\mathbf{x} \in \mathbb{R}^n$ and
- (2) $\mathcal{L}V(\mathbf{x}) \leq -c_3 \|\mathbf{x} - \mathbf{x}^*\|^2 + \gamma'$ for all $\mathbf{x} \in \mathbb{R}^n$.

Then, the state $\mathbf{x}^*$ is *globally practically uniformly exponentially stable in the mean square sense* for (1), with parameters $\alpha = \frac{c_2}{c_1}$, $\beta = \frac{c_3}{c_2}$ and $\gamma = \frac{\gamma'}{c_3}$.

III. MAIN RESULTS

In this section, we apply the previously established SDE framework to bridge the gap between idealized communication models often assumed in distributed optimization and the sparse, asynchronous reality of networked systems. Our goal is to design a distributed gradient descent algorithm where communication constraints are integral to the dynamics.

A. Channel Dynamics Model

To realise distributed algorithms within our SDE framework, we model the interaction between any two agents ν_j and ν_i connected by an edge $(j, i) \in \mathcal{E}$. As depicted in Figure 2, each agent ν_i maintains its own state $\mathbf{x}_i$ governed by continuous-time dynamics. Information exchange between agents, however,




Fig. 2: Illustration of two nodes within the communication graph and their directed communication channels: solid lines represent continuous signals, while dashed lines represent discrete communication events. Channel and node dynamics are continuous-time dynamics.

occurs only through discrete, stochastic communication events modeled as Poisson jumps. The continuous-time evolution of the agent dynamics can be interpreted as a time scale separation in the sense that sampling rates of the agents are much faster than the ones over the communication channel.

Specifically, for each directed communication link $(j, i) \in \mathcal{E}$, agent ν_i maintains a local channel state $\mathbf{z}_{(j,i)}(t, \omega) \in \mathbb{R}^{d_j}$. This state represents the information agent ν_i possesses about agent ν_j 's state $\mathbf{x}_j$. The crucial aspect is that $\mathbf{z}_{(j,i)}$ is only updated sporadically when a communication event (a Poisson jump) occurs on the channel (j, i) . Between these events, the channel state might evolve continuously (e. g., decay) or remain constant. The general dynamics for the channel state $\mathbf{z}_{(j,i)}$ and its output $\mathbf{y}_{(j,i)}$ (which is then used by agent ν_i) are given by the SDE

$$d\mathbf{z}_{(j,i)}(t, \omega) = f(\mathbf{z}_{(j,i)}(t, \omega)) dt \quad (4a)$$

$$+ g(\mathbf{x}_j(t, \omega), \mathbf{z}_{(j,i)}(t, \omega)) dN_{(j,i)}(t, \omega),$$

$$\mathbf{y}_{(j,i)}(t, \omega) = h(\mathbf{z}_{(j,i)}(t, \omega)). \quad (4b)$$

Here, $dN_{(j,i)}$ is the Poisson process modelling communication events from ν_j to ν_i with average rate $\lambda_{(j,i)}$. The map $f : \mathbb{R}^{d_{(j,i)}} \rightarrow \mathbb{R}^{d_{(j,i)}}$ describes the continuous




Fig. 3: Sample paths with different channel dynamics, showing the continuous evolution of $(\mathbf{x}_1(t, \omega), \dots)$ and $(\mathbf{z}_{(1,2)}(t, \omega), \dots)$ interrupted by discrete communication events $dN_{(1,2)}(t, \omega)$.

evolution of the channel state between communications, while $g : \mathbb{R}^{d_j} \times \mathbb{R}^{d_{(j,i)}} \rightarrow \mathbb{R}^{d_{(j,i)}}$ defines how the state $\mathbf{x}_j$ of the sender updates the channel state $\mathbf{z}_{(j,i)}$ at the instant of a communication event. The output map $h : \mathbb{R}^{d_{(j,i)}} \rightarrow \mathbb{R}^{m_{(j,i)}}$ determines what information $\mathbf{y}_{(j,i)}$ is actually available to the receiving agent ν_i based on the channel state.

This model explicitly separates the continuous internal dynamics of each agent from the discrete, event-driven communication between them. It provides flexibility to represent various channel behaviours, including potentially unstable dynamics between updates, i. e., f could represent an unstable system, allowing for modelling scenarios like accumulating errors or decaying information quality.

We explore two representative effects of Poisson jumps on the channel output $\mathbf{y}_{(j,i)}$, illustrating the modelling approach through practical scenarios contrasted in Figure 3

1) *Stochastic Ideal Sample-and-Hold*: In environmental monitoring networks, sensors transmit measurements like temperature or humidity to a central hub, which must retain the latest vector of readings exactly until a new message arrives. This requires a channel that captures ν_j 's multi-dimensional state instantly and holds it without degradation, akin to an ideal sample-and-hold process in signal processing.

Example 1 (Ideal Sample-and-Hold Channel) Consider a channel syncing ν_i 's data with ν_j 's state at message arrivals. The dynamics are

$$d\mathbf{z}_{(j,i)} = (\mathbf{x}_j - \mathbf{z}_{(j,i)}) dN_{(j,i)}, \quad (5a)$$

$$\mathbf{y}_{(j,i)} = \mathbf{z}_{(j,i)}, \quad (5b)$$

where $\mathbf{z}_{(j,i)} \in \mathbb{R}^{d_j}$. At each jump, $\mathbf{z}_{(j,i)}$ resets to $\mathbf{x}_j(t^-)$, delivering ν_j 's state to ν_i without loss.

As seen in Figure 3a, the channel output $\mathbf{z}_{(j,i)}$ perfectly holds the value of $\mathbf{x}_j$ sampled at the last jump time, until the next jump occurs, while $\mathbf{x}_j$ proceeds to evolve in continuous-time.

2) *Leaky Integrator Channel*: In low-cost sensor hardware, channels aim to store values but suffer from gradual loss, as in a switched RC circuit where a capacitor leaks charge through non-ideal resistive effects, as shown in Figure 4. More ideal signal holding behavior is possible, but comes at great cost in complex refresh circuits or the addition of non-volatile memories. The same effect can be observed in neurons, where activation abates over time. This smooths incoming messages over time, resembling a faulty sample-and-hold process.

Example 2 (Leaky Integrator Channel) Consider a channel mimicking an RC circuit. The dynamics are

$$d\mathbf{z}_{(j,i)} = a_{(j,i)} \mathbf{z}_{(j,i)} dt + (\mathbf{x}_j - \mathbf{z}_{(j,i)}) dN_{(j,i)}, \quad (6a)$$




Fig. 4: Switched RC circuit modelling a leaky integrator channel, where input x_j charges the capacitor at Poisson jumps, and output $z_{(j,i)}$ leaks through R .

$$\mathbf{y}_{(j,i)} = \mathbf{z}_{(j,i)}, \quad (6b)$$

where $\mathbf{z}_{(j,i)} \in \mathbb{R}^{d_j}$. Here, $a_{(j,i)} \in \mathbb{R}$ is the decay (or potentially amplification, if $a_{(j,i)} > 0$) rate. Figure 3b shows how the channel state $\mathbf{z}_{(j,i)}$ jumps to match $\mathbf{x}_j$ but then decays according to e^{at} until the next communication event. —

Naturally, a multitude of other forms of channel dynamics are possible, including stochastic delay channels, synaptic integration channels, consensus channels or observer channels, but we believe the two exemplified versions are the most common and serve as a good presentational foundation for the remainder of the paper.

These examples demonstrate how the SDE framework allows us to model the interplay between continuous node computations and discrete communication events, capturing channel imperfections like information decay or even instability. The key question, addressed next, is how algorithms designed using this communication model maintain stability and performance.

B. Distributed Gradient Descent

This section applies the communication model introduced in Section III-A to a well-known problem in optimization, namely distributed gradient descent for quadratic cost functions. In doing so, we start with a formal definition of the problem, then motivate the construction of distributed optimization dynamics governed by our novel Poisson jump communication channel model and conclude with a presentation of the main theoretical contributions, each accompanied by discussions concerning their broader implications.

Our exploration begins with the formal definition of the reference problem, which serves as the foundational example for illustrating our proposed model.

Definition 2 (Quadratic Optimization Problem) Consider the nominal unconstrained quadratic optimization problem

$$\min_{\mathbf{y} \in \mathbb{R}^d} J(\mathbf{y}) := \frac{1}{2} \mathbf{y}^T \mathbf{Q} \mathbf{y} + \mathbf{q}^T \mathbf{y}, \quad (7)$$

where $\mathbf{Q} = \mathbf{Q}^T \succ 0$ is symmetric and positive definite and $\mathbf{q} \in \mathbb{R}^d$. The unique optimal solution is given by $\mathbf{y}^* = -\mathbf{Q}^{-1} \mathbf{q}$. —

A well-known method for asymptotically solving this quadratic optimization problem is to formulate it as a continuous-time dynamic system, in particular as the gradient descent flow.

Definition 3 (Gradient Descent Flow) Associated with (7), the gradient descent flow is

$$d\mathbf{y} = -(\mathbf{Q}\mathbf{y} + \mathbf{q}) dt. \quad (8)$$

In anticipation of the parallels to be drawn with a stochastic setup, we have elected to present the gradient descent flow in a differential formulation.

In view of our desire to split this nominal problem into smaller sub-problems, we can partition $\mathbf{y}$ into n sub-vectors $\mathbf{y}_p$, and partition $\mathbf{Q}$ and $\mathbf{q}$ accordingly, resulting in

$$\mathbf{Q} = \begin{pmatrix} \mathbf{Q}_{11} & \cdots & \mathbf{Q}_{1n} \\ \vdots & \ddots & \vdots \\ \mathbf{Q}_{n1} & \cdots & \mathbf{Q}_{nn} \end{pmatrix}, \quad \mathbf{q} = \begin{pmatrix} \mathbf{q}_1 \\ \vdots \\ \mathbf{q}_n \end{pmatrix}.$$

The resulting gradient of the global objective function J with respect to an individual sub-vector $\mathbf{y}_i$ can then be precisely expressed as

$$\nabla_{\mathbf{y}_i} J(\mathbf{y}) = \mathbf{Q}_{ii} \mathbf{y}_i + \sum_{j \neq i} \mathbf{Q}_{ij} \mathbf{y}_j + \mathbf{q}_i. \quad (9)$$

Consequently, the nominal gradient descent flow (8), when considered from the vantage point of each partitioned component $\mathbf{y}_i$, would ideally evolve according to

$$d\mathbf{y}_i = -(\mathbf{Q}_{ii} \mathbf{y}_i + \sum_{j \neq i} \mathbf{Q}_{ij} \mathbf{y}_j + \mathbf{q}_i) dt. \quad (10)$$

This formulation can already be thought of as a distributed algorithm, however at each time instant the vectorfield describing the evolution of $\mathbf{y}_i$ requires knowledge of the neighbour sub-vectors $\mathbf{y}_j$. In a distributed system, for an agent ν_i , the sub-vectors $\mathbf{y}_j$ for $j \neq i$ will generally be not available, so our idea is to replace them with the channel output $\mathbf{z}_{(j,i)}$ of our proposed channel model instead. Jointly with the channel dynamics, this gives rise to our proposed distributed gradient descent flow with Poisson-distributed communication events.

Definition 4 Associated with (7), the distributed gradient descent flow with Poisson-distributed communication events for agent ν_i with $i \in [1, n] \cap \mathbb{N}$ and $j \in [0, n] \cap \mathbb{N} \setminus \{i\}$, is

$$d\mathbf{x}_i(t) = -(\mathbf{Q}_{ii} \mathbf{x}_i(t) + \sum_{j \neq i} \mathbf{Q}_{ij} \mathbf{z}_{(j,i)}(t) + \mathbf{q}_i) dt, \quad (11a)$$

$$d\mathbf{z}_{(j,i)}(t) = a_{(j,i)}(t) \mathbf{z}_{(j,i)}(t) dt + (\mathbf{x}_j(t) - \mathbf{z}_{(j,i)}(t)) dN_{(j,i)}(t), \quad (11b)$$

where $\mathbf{Q}_{ij} \in \mathbb{R}^{d_i \times d_j}$, $\mathbf{q}_i \in \mathbb{R}^{d_i}$, $a_{(j,i)}(t) \in \mathbb{R}$, and $N_{(j,i)}$ are independent Poisson processes with rates $\lambda_{(j,i)}$. —

System (11) highlights the application of our modelling technique. Firstly, channel dynamics (11b) for communication events from agents ν_j are augmented to each agent ν_i in an existing, well known algorithm, and then the algorithm is adapted to use the local copies from the channel dynamics where applicable (11a).

For easier analysis, we switch to an equivalent representation using error dynamics relative to the unknown optimal solution $\mathbf{y}^*$.

Lemma 2 (Error Dynamics) Consider the distributed gradient descent flow (11), associated with the quadratic optimization problem (7). Let the error be $\tilde{\mathbf{x}} := \mathbf{x} - \mathbf{y}^* \in \mathbb{R}^d$. Furthermore, let the copy errors be $\tilde{\mathbf{z}}_{(j,i)} := \mathbf{z}_{(j,i)} - \mathbf{y}_j^* \in \mathbb{R}^{d_j}$ for agent ν_i 's copy of agent ν_j 's state $\mathbf{x}_j$, with $i \in [1, n] \cap \mathbb{N}$ and $j \in [0, n] \cap \mathbb{N} \setminus \{i\}$. The resultant error dynamics are

$$d\tilde{\mathbf{x}}_i(t) = -\left(\mathbf{Q}_{ii} \tilde{\mathbf{x}}_i(t) + \sum_{j \neq i} \mathbf{Q}_{ij} \tilde{\mathbf{z}}_{(j,i)}(t)\right) dt, \quad (12a)$$

$$d\tilde{\mathbf{z}}_{(j,i)}(t) = \left(a_{(j,i)}(t) \tilde{\mathbf{z}}_{(j,i)}(t) + a_{(j,i)}(t) \mathbf{y}_j^*\right) dt + (\tilde{\mathbf{x}}_j(t) - \tilde{\mathbf{z}}_{(j,i)}(t)) dN_{(i,p)} \quad (12b)$$

The error system will be used for analysis only, implementation must be done using (11), which we will refer to for the upcoming stability result. In particular, we will be considering the stability of a state $(\mathbf{x}^*, \mathbf{z}^*) = (\mathbf{y}^*, \text{vec}(\mathbf{y}_{-1}^*, \dots, \mathbf{y}_{-n}^*))$, as (11) does not exhibit a trivial steady-state solution in general. In the language of Definition 1, we can state the following theorem.

Theorem 1 Consider the distributed gradient descent system (11). Assume that for each $a_{(j,i)}(t)$, there exists a bound a_{ji} , such that $|a_{(j,i)}(t)| \leq a_{ji}$ for all $t \in [t_0, \infty)$. Then,

- (a) for every $\gamma \in (0, \infty)$, there exists a constant communication rate $\lambda_s \in [0, \infty)$, such that for all $\lambda_{(j,i)} \in (\lambda_s, \infty)$, the state $(\mathbf{x}^*, \mathbf{z}^*)$ is *globally practically uniformly exponentially stable in the mean-square sense* for (11) with ultimate bound γ ;
- (b) for every $\gamma \in (0, \infty)$ and $\beta \in (0, 2\mu_{\min}(\mathbf{Q}))$, there exists a constant communication rate $\lambda_d \in [\lambda_s, \infty)$ such that for all choices $\lambda_{(j,i)} \in (\lambda_d, \infty)$, the stability properties from (a) hold and the exponential convergence rate in (3) is β .

Remark 1 The practical stability bound γ in Definition 1 is achieved by selecting positive weighting parameters $(\rho_j)_{j=1..n}$ in the proof (see Step 2). These parameters must be chosen to satisfy $\gamma' \leq c_3\gamma$, where c_3 is the resulting convergence rate (e.g., $c_3 = \beta$ for statement (b) of Theorem 1) and $\gamma' = \sum_{i=1}^n \sum_{j \neq i} \frac{1}{\rho_j} \|\mathbf{y}_j^*\|^2$. This reveals a fundamental trade-off: Achieving a smaller (more accurate) ultimate bound γ requires choosing larger weighting parameters ρ_j . Larger ρ_j generally lead to more restrictive bounds in the stability analysis (Step 5), which in turn demand higher sufficient communication rates λ_s and λ_d . The size of the ultimate bound, characterized by γ , also scales with the number n of interfering channels and the magnitude of the optimal solution $\mathbf{y}^*$.

Remark 2 Sufficient choices for the stabilizing rate λ_s and the convergence-guaranteeing rate λ_d can be analytically obtained. Their derivation, which depends on bounds of the system matrices and the chosen weighting parameters $(\rho_j)_{j=1, \dots, n}$, is detailed in the proof of this theorem, see (23).

The proof of this theorem is deferred to the appendix in favour of a more streamlined presentation. Observe, that the theorem states the convergence rate and practical stability for the variance of the solutions to the error system. However, second moment stability entails first moment stability [21], i.e. we also have

$$\mathbb{E}[\|\mathbf{s}\|] \leq \sqrt{c_3} \mathbb{E}[\|\mathbf{s}^0\|] e^{-\frac{1}{2}\beta(t-t_0)} + \sqrt{\gamma}$$

by application of Jensen's inequality and by subadditivity of $\sqrt{\cdot}$.

The first part of the theorem states sufficient conditions on the communication channels for achieving stability. The second part notably relates the exponential convergence rate to the convergence rate of the nominal problem and it can be observed, that for sufficiently high communication rates the convergence in expectation can be designed to achieve any rate up to the nominal convergence rate from the nominal gradient descent flow in Definition 3. The overall convergence rates will be limited by $\mu_{\min}(\mathbf{Q})$, for the nominal as well as for the distributed gradient descent.

If the communication channel lack drifts, however, we can conveniently state an even tighter result.

Corollary 1 Consider the distributed gradient descent system (11) with zero channel drift rates (i.e., $a_{(j,i)}(t) \equiv 0$ for all communication channels (j, i)). Then, in Theorem 1 the qualifier *practical* may be removed, i.e. $\gamma = 0$.

Unfortunately, the sufficient conditions on λ_s and λ_d are only computable by inverting $\mathbf{Q}$, solving large Linear Matrix Inequalities or assuming a particular beneficial property like diagonal dominance of $\mathbf{Q}$, which are an obstacle in a completely distributed setup, where ideally each agent should be able to determine its own communication rates independently from its neighbours. As such, our result remains mostly an existence result, but might be usable nonetheless, as manual tuning of rates or adaptation guarantees stability for sufficiently high rates.

IV. SIMULATION

To illustrate our theoretical findings from the previous section, we apply (11) to an academic example. We consider a quadratic optimization problem of form (7) with

$$\mathbf{Q} = \begin{bmatrix} 4 & 2 & 1 & 1 & 0 & 2 \\ 2 & 5 & 2 & 0 & 2 & 1 \\ 1 & 2 & 6 & 3 & 1 & 0 \\ 1 & 0 & 3 & 4 & 2 & 1 \\ 0 & 2 & 1 & 2 & 5 & 2 \\ 2 & 1 & 0 & 1 & 2 & 4 \end{bmatrix} \quad \text{and} \quad \mathbf{q} = \begin{bmatrix} -9 \\ -15 \\ -22 \\ -12 \\ -10 \\ -5 \end{bmatrix}$$

and formulate our proposed distributed gradient descent flow with Poisson-distributed communication rates for $n = 3$ heterogeneous agents with respective state dimensions $d_1 = 3$, $d_2 = 2$, $d_3 = 1$. We deliberately have chosen a symmetric positive definite matrix $\mathbf{Q}$, which is not diagonally dominant to highlight this is not a required property. The nominal convergence rate is characterized by $\mu_{\min}(\mathbf{Q}) \approx 0.31$.

We simulate our proposed algorithm in different scenarios for different choices of communication rates as well as for channels with and without drift. For the continuous part of the SDE we used an *ode1*-solver with step-size $h = 0.01$. Throughout the simulations, we have opted to select all communication and drift rates $\lambda = \lambda_i$ equal for a more compact presentation in the paper.

Experiment 1 Let $\mathbf{a} = \mathbf{0}$. In view of Theorem 1, we obtain numerically the sufficient communication rate $\lambda_s \approx 26.56$. Let $\lambda_1 = 10$, $\lambda_2 = 27$ and $\lambda_3 = 50$.

The simulation results are depicted in Figure 5, which shows a comparison of the evolution of the Lyapunov-function (13) for different communication rates alongside the the nominal convergence bound $\frac{3}{2}e^{-2\mu_{\min}(\mathbf{Q})t}$. A single selected sample path $V(\mathbf{s}(\cdot, \omega_1))$ is shown, which showcases, that individual sample paths of V need not be decreasing at all times in this stochastic setting, unlike for a deterministic system.

Taking the average $\bar{V}(t) := \frac{1}{N} \sum_{k=1}^N V(\mathbf{s}(t, \omega_k))$ over $N = 100$ sample paths of $V(\mathbf{s}(t, \omega_k))$ however illustrates the statement of Theorem 1 and highlights the exponential stability in the mean-square sense. A 95-percentile confidence interval for $\bar{V}$ is depicted alongside, to illustrate the spread of trajectories and the exponential decline of the sample path variance.




Fig. 5: Comparison of Lyapunov-function evolutions with zero channel drift rates ($\mathbf{a} = \mathbf{0}$). Depicted are $\bar{V}(t)$ for different communication rates: ($\lambda_1 = 10$, —), ($\lambda_2 = 27$, —) and ($\lambda_3 = 50$, —) alongside their respective confidence intervals, as well as a reference exponential function with decay rate ($2\mu_{\min}(Q)$, - - -) and a single sample path (- - -).




Fig. 6: Comparison of Lyapunov-function evolutions with unstable channels ($\mathbf{a}_{ji} = 1$ for $j \neq i$). Depicted are $\bar{V}(t)$ for different communication rates ($\lambda_1 = 26$, —), ($\lambda_2 = 51$, —).

Regarding the communication rates λ , the following can be observed:

- $\lambda = \lambda_1 < \lambda_s$ Even with communication rates smaller than λ_s , we observe convergence at ultimately nearly identical rates to the optimal solution in some cases, which was not part of the statement in Theorem 1.
- $\lambda \in \{\lambda_2, \lambda_3\}, \lambda_r < \lambda_3 < \lambda_4$ We observe, that increasing the communication rates beyond λ_s only yields insignificant improvements for the convergence rate of the Lyapunov function, unlike discussed.

Experiment 2 Let $\mathbf{a}_{ip} = 1$ for each $i \neq p$ and $\mathbf{a}_{ii} = 0$. In view of Theorem 1, we obtain numerically the sufficient communication rate $\lambda_s \approx 50.57$. Let $\lambda_1 = 26$ and $\lambda_2 = 51$.

The simulation results are depicted in Figure 6, which shows a comparison of the evolution of the Lyapunov-function (13) for different communication rates. Once again, the averages $\bar{V}(t)$ and 95-percentile confidence interval for $\bar{V}(t)$ are depicted. Regarding the communication rates λ , the following can be observed:

- $\lambda = \lambda_1 < \lambda_s$ Even with communication rates smaller than λ_s , we observe practical convergence, which was not part of the statement in Theorem 1.
- $\lambda = \lambda_2 > \lambda_s$ Increasing the communication rates means essentially selecting smaller ultimate bound γ (or equivalently: higher scaling weights ρ) and renders the disk of practical convergence smaller, as stated by Remark 1.

V. DISCUSSION, CONCLUSION & OUTLOOK

In this paper, we have presented a novel framework for designing communication-aware distributed optimization algorithms. The core of our approach, motivated by networked systems and modern hardware, is to model the communication channel itself as a dynamic, stochastic process using SDEs

driven by Poisson Jumps. This fundamentally shifts the design paradigm: instead of treating communication scarcity as an afterthought imposed on an existing algorithm—as is common in many event-triggered schemes—our framework embeds these channel constraints into the mathematical flow from its inception. This allows us to take nominal continuous-time dynamics (representing ideal, fully-connected systems) and rigorously analyze their stability and performance under sparse, asynchronous, and stochastic communication events, such as exponentially distributed sampling times. Unlike state-dependent event-triggering, this Poisson-driven approach does not require agents to continuously monitor local state changes to make communication decisions, which significantly reduces the computational and energy overhead associated with triggering logic. We have illustrated the process of turning a continuous-time dynamic system (the gradient descent flow) into a network system with constant communication (just a matter of different interpretation of the gradient descent flow), then applying our proposed channel dynamics model to obtain the distributed gradient descent flow with Poisson-distributed communication events. Using our modelling technique, we were then able to state stability guarantees depending on sufficiently high communication rates in two increasingly meaningful ways: Firstly, we show existence of sufficient stabilizing communication rates, and secondly, we show the existence of sufficient communication rates to achieve a desired convergence performance up to the nominal gradient descent flow. While it serves as an example to illustrate our modelling technique, the proposed distributed gradient descent flow with Poisson-distributed communication events is a usable algorithm on its own. Even though the chosen illustrative example (the distributed gradient descent flow) requires a complete communication graph due to its reliance on global state aggregation, this is not a limitation of the SDE framework itself. The framework's core strength lies in its ability to model and analyze complex (deterministic and stochastic) channel dynamics, including time-varying or unstable channels or sparse random communication patterns over directed graphs. Applying the Poisson-driven SDE approach to consensus-based algorithms designed for sparse topologies (e.g., DGD or ADMM) presents one natural direction for future work.

Another future research direction is to apply the approach to algorithms that are natively designed for sparse communication graphs, such as consensus-based DGD or ADMM. This would directly connect our theoretical framework to the sparsely-connected hardware architectures that motivated this work. Further, the framework's flexibility invites extensions to more complex problem classes, including constrained quadratic programming and the integration of state observers to handle partial information.

Moreover, a promising direction for future work involves exploring accelerated convergence within this distributed SDE framework. Leveraging techniques from convex synthesis for optimization algorithms for centralised systems [23], could enable the co-design of both the algorithm dynamics and the communication strategy. This might lead to characterising minimal communication rates to achieve performance exceeding the nominal gradient flow, even under sparse, stochastic communication constraints. Finally, our focus on continuous-time dynamics positions could be transferred to discrete-time models with jump dynamics, such as [24]. Investigating

the connections and potential translations between these continuous and discrete-time stochastic frameworks represents a fertile ground for further future theoretical exploration, promising a deeper understanding of distributed optimization.

REFERENCES

- [1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *Artificial Intelligence and Statistics*. PMLR, 2017, pp. 1273–1282.
- [2] D. K. Molzahn, F. Dörfler, H. Sandberg, S. H. Low, S. Chakrabarti, R. Baldick, and J. Lavaei, “A survey of distributed optimization and control algorithms for electric power systems,” *IEEE Transactions on Smart Grid*, vol. 8, no. 6, pp. 2941–2962, 2017.
- [3] T. Yang, X. Yi, J. Wu, Y. Yuan, D. Wu, Z. Meng, Y. Hong, H. Wang, Z. Lin, and K. H. Johansson, “A survey of distributed optimization,” *Annual Reviews in Control*, vol. 47, pp. 278–305, 2019.
- [4] A. Nedic and A. Ozdaglar, “Distributed subgradient methods for multi-agent optimization,” *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 48–61, 2009.
- [5] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [6] W. Heemels, K. H. Johansson, and P. Tabuada, “An introduction to event-triggered and self-triggered control,” in *51st IEEE Conference on Decision and Control 2012, Hawaii, USA*. IEEE conference proceedings, 2012, pp. 3270–3285.
- [7] P. Tabuada, “Event-triggered real-time scheduling of stabilizing control tasks,” *IEEE Transactions on Automatic Control*, vol. 52, no. 9, pp. 1680–1685, 2007.
- [8] G. Notarstefano, I. Notarnicola, A. Camisa *et al.*, “Distributed optimization for smart cyber-physical networks,” *Foundations and Trends® in Systems and Control*, vol. 7, no. 3, pp. 253–383, 2019.
- [9] G. Carnevale, I. Notarnicola, L. Marconi, and G. Notarstefano, “Triggered gradient tracking for asynchronous distributed optimization,” *Automatica*, vol. 147, p. 110726, 2023.
- [10] AMD, “4th gen amd epyc processor architecture,” Advanced Micro Devices (AMD), White Paper, Nov 2022. [Online]. Available: <https://www.amd.com/content/dam/amd/en/documents/products/epyc/4th-gen-epyc-processor-architecture-white-paper.pdf>
- [11] C. Mead, “Neuromorphic electronic systems,” *Proceedings of the IEEE*, vol. 78, no. 10, pp. 1629–1636, 2002.
- [12] M. Davies, N. Srinivasa, T.-H. Lin, G. Chinya, Y. Cao, S. H. Choday, G. Dimou, P. Joshi, N. Imam, S. Jain *et al.*, “Loihi: A neuromorphic manycore processor with on-chip learning,” *IEEE Micro*, vol. 38, no. 1, pp. 82–99, 2018.
- [13] N. Leroux, P.-P. Manea, C. Sudarshan, J. Finkbeiner, S. Siegel, J. P. Strachan, and E. Neftci, “Analog in-memory computing attention mechanism for fast and energy-efficient large language models,” *Nature Computational Science*, pp. 1–12, 2025.
- [14] G. H. Hutchinson, E. Sifferman, T. Bhattacharya, D. Kwon, and D. B. Strukov, “Fpia: Field-programmable ising arrays with in-memory computing,” in *Proceedings of the 29th ACM/IEEE International Symposium on Low Power Electronics and Design*, ser. ISLPED ’24. New York, NY, USA: Association for Computing Machinery, 2024, p. 1–6. [Online]. Available: <https://doi.org/10.1145/3665314.3670851>
- [15] R. Brockett, “Stochastic control,” Apr. 2009, lecture Notes, Version of April 3, 2009.
- [16] X.-M. Zhang, Q.-L. Han, X. Ge, D. Ding, L. Ding, D. Yue, and C. Peng, “Networked control systems: A survey of trends and techniques,” *IEEE/CAA Journal of Automatica Sinica*, vol. 7, no. 1, pp. 1–17, 2019.
- [17] R. Brockett and G. Blankenship, “A representation theorem for linear differential equations with Markovian coefficients,” in *Proceedings for the 15th Allerton Conference on Circuits and Systems Theory*, 1977, pp. 671–679.
- [18] M. H. A. Davis, *Markov models and optimization*, Springer reprint, *Monographs on Statistics and Applied Probability*. Chapman and Hall, 1993.
- [19] S. Summers and J. Lygeros, “Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem,” *Automatica*, vol. 46, no. 12, pp. 1951–1961, 2010.
- [20] F. Farokhi and K. H. Johansson, “Stochastic sensor scheduling for networked control systems,” *IEEE Transactions on Automatic Control*, vol. 59, no. 5, pp. 1147–1162, 2014.
- [21] X. Mao, *Stochastic differential equations and applications*. Elsevier, 2007.
- [22] E. D. Sontag, “Input to state stability: Basic concepts and results,” in *Nonlinear and Optimal Control Theory: Lectures given at the CIME Summer School held in Cetraro, Italy June 19–29, 2004*. Springer, 2008, pp. 163–220.
- [23] C. Scherer and C. Ebenbauer, “Convex synthesis of accelerated gradient algorithms,” *SIAM Journal on Control and Optimization*, vol. 59, no. 6, pp. 4615–4645, 2021.
- [24] C. Ebenbauer and R. Suttner, “A piecewise deterministic discrete-time model with simultaneous random jumps,” *Systems & Control Letters*, vol. 203, p. 106141, 2025.
- [25] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 2012.

APPENDIX A: LEMMATA

Lemma 3 Let M be a symmetric block matrix of the form

$$M = \begin{bmatrix} M_{11} & M_{12} \\ M_{12}^T & \Lambda + R \end{bmatrix}$$

where M_{11} is symmetric and positive definite, R is symmetric and $\Lambda = \text{diag}(\lambda_k)$. The following two implications hold:

- (1) Let $K := M_{12}^T M_{11}^{-1} M_{12}$. If all λ_k satisfy

$$\lambda_k > \lambda_s := -\mu_{\min}(R - K),$$

then the matrix M is positive definite.

- (2) Let $\mu \in (0, \mu_{\min}(M_{11}))$ and $\tilde{K}_\mu := M_{12}^T(M_{11} - \mu I)^{-1} M_{12}$. If all λ_k satisfy

$$\lambda_k \geq \lambda_d := \mu - \mu_{\min}(R - \tilde{K}_\mu),$$

then $M \succeq \mu I$.

APPENDIX B: PROOFS

Proof (of Lemma 3)

(1) A symmetric matrix M is positive definite if and only if its principal submatrix M_{11} is positive definite and its Schur complement $S := \Lambda + R - K$ is positive definite [25, p. 495]. According to the assumptions of this lemma, M_{11} is positive definite. Thus, we need to show that S is positive definite, which is equivalent to its smallest eigenvalue $\mu_{\min}(S)$ being strictly positive. We lower-bound $\mu_{\min}(S)$ via

$$\mu_{\min}(S) = \mu_{\min}(\Lambda + (R - K)) \geq \mu_{\min}(\Lambda) + \mu_{\min}(R - K)$$

using Weyl’s inequality for eigenvalues of sums of symmetric matrices [25, Corollary 4.3.15].

Since $\Lambda = \text{diag}(\lambda_k)$, its smallest eigenvalue is $\mu_{\min}(\Lambda) = \min_k(\lambda_k)$. Therefore, for S to be positive definite, we require

$$\min_k(\lambda_k) + \mu_{\min}(R - K) > 0.$$

Rearranging yields the first statement of this lemma.

(2) The condition $M \succeq \mu I$ is equivalent to showing $M' := M - \mu I \succeq 0$. Since $\mu < \mu_{\min}(M_{11})$, the block $M'_{11} \succ 0$ and subtracting the diagonal matrix μI inherits the symmetry from M to M' . Applying the first statement of this lemma to M' and considering positive semi-definiteness of S yields the second statement of this lemma. $\square$

Proof (of Theorem 1) This proof aims to show that the system is globally practically uniformly exponentially stable in the mean-square sense by applying the Lyapunov theorem from Lemma 1.

Let $\gamma \in (0, \infty)$ be given. For the second statement, let $\beta \in (0, 2\mu_{\min}(\mathbf{Q}))$ also be given.

(Step 1) For notational convenience, define firstly $\mathbf{s}^\top := [\tilde{\mathbf{x}}^\top \mathbf{e}^\top]$, where $\mathbf{e}$ is the stacked vector of copy errors, with $\mathbf{e}_{(j,i)} := \tilde{\mathbf{x}}_j - \tilde{\mathbf{z}}_{(j,i)} \in \mathbb{R}^{d_j}$. As a candidate Lyapunov function, we choose the quadratic form

$$V(\mathbf{s}) := \|\mathbf{s}\|^2 = \sum_{i=1}^n \|\tilde{\mathbf{x}}_i\|^2 + \sum_{i=1}^n \sum_{j \neq i} \|\mathbf{e}_{(j,i)}\|^2. \quad (13)$$

This function satisfies the first condition of Lemma 1 with $c_1 = 1$ and $c_2 = 1$. According to Lemma 1, we need to show that there exist constants $c_3 > 0$ and $\gamma' > 0$ such that $\mathcal{L}V(\mathbf{s}) \leq -c_3 \|\mathbf{s}\|^2 + \gamma'$, where the resulting stability parameters satisfy $\beta = \frac{c_3}{c_2} = c_3$ and $\gamma' = c_3 \gamma$.

(Step 2) We directly calculate the infinitesimal generator

$$\mathcal{L}V(\mathbf{s}) \quad (14)$$

$$\begin{aligned} & \stackrel{(a)}{=} \sum_{i=1}^n \left(\frac{\partial V}{\partial \tilde{\mathbf{x}}_i} \right)^\top \left(-\mathbf{Q}_{ii} \tilde{\mathbf{x}}_i(t) - \sum_{j \neq i} \mathbf{Q}_{ij} \tilde{\mathbf{z}}_{(j,i)}(t) \right) \\ & + \sum_{i=1}^n \sum_{j \neq i} \left(\frac{\partial V}{\partial \tilde{\mathbf{z}}_{(j,i)}} \right)^\top \left(a_{(j,i)}(t) \tilde{\mathbf{z}}_{(j,i)}(t) + a_{(j,i)}(t) \mathbf{y}_j^* \right) \\ & + \sum_{i=1}^n \sum_{j \neq i} \lambda_{(j,i)} (V(\mathbf{s} + \Delta \mathbf{s}) - V(\mathbf{s})) \\ & \stackrel{(b)}{=} 2 \sum_{i=1}^n \left(\tilde{\mathbf{x}}_i + \sum_{j \neq i} (\tilde{\mathbf{x}}_j - \tilde{\mathbf{z}}_{(j,i)}) \right)^\top \left(-\mathbf{Q}_{ii} \tilde{\mathbf{x}}_i - \sum_{j \neq i} \mathbf{Q}_{ij} \tilde{\mathbf{z}}_{(j,i)} \right) \\ & - 2 \sum_{i=1}^n \sum_{j \neq i} (\tilde{\mathbf{x}}_j - \tilde{\mathbf{z}}_{(j,i)})^\top (a_{(j,i)}(t) \tilde{\mathbf{z}}_{(j,i)} + a_{(j,i)}(t) \mathbf{y}_j^*) \\ & - \sum_{i=1}^n \sum_{j \neq i} \lambda_{(j,i)} (\tilde{\mathbf{x}}_j - \tilde{\mathbf{z}}_{(j,i)})^\top (\tilde{\mathbf{x}}_j - \tilde{\mathbf{z}}_{(j,i)}) \\ & \stackrel{(c)}{=} -2 \sum_{i=1}^n \tilde{\mathbf{x}}_i^\top \mathbf{Q}_{ii} \tilde{\mathbf{x}}_i \\ & - 2 \sum_{i=1}^n \sum_{j \neq i} \tilde{\mathbf{x}}_i^\top \mathbf{Q}_{ij} (\tilde{\mathbf{x}}_j - (\tilde{\mathbf{x}}_j - \tilde{\mathbf{z}}_{(j,i)})) \\ & - 2 \sum_{i=1}^n \sum_{j \neq i} (\tilde{\mathbf{x}}_i - \tilde{\mathbf{z}}_{(i,j)})^\top \mathbf{Q}_{ii} \tilde{\mathbf{x}}_i \\ & - 2 \sum_{i=1}^n \sum_{j \neq i} \sum_{k \neq i} (\tilde{\mathbf{x}}_i - \tilde{\mathbf{z}}_{(i,j)})^\top \mathbf{Q}_{ik} (\tilde{\mathbf{x}}_k - (\tilde{\mathbf{x}}_k - \tilde{\mathbf{z}}_{(k,i)})) \\ & - 2 \sum_{i=1}^n \sum_{j \neq i} (\tilde{\mathbf{x}}_j - \tilde{\mathbf{z}}_{(j,i)})^\top a_{(j,i)}(t) (\tilde{\mathbf{x}}_j - (\tilde{\mathbf{x}}_j - \tilde{\mathbf{z}}_{(j,i)})) \\ & - 2 \sum_{i=1}^n \sum_{j \neq i} (\tilde{\mathbf{x}}_j - \tilde{\mathbf{z}}_{(j,i)})^\top a_{(j,i)}(t) \mathbf{y}_j^* \\ & - \sum_{i=1}^n \sum_{j \neq i} \lambda_{(j,i)} (\tilde{\mathbf{x}}_j - \tilde{\mathbf{z}}_{(j,i)})^\top (\tilde{\mathbf{x}}_j - \tilde{\mathbf{z}}_{(j,i)}) \end{aligned}$$

$$\begin{aligned} & \stackrel{(d)}{=} -2 \sum_{i=1}^n \tilde{\mathbf{x}}_i^\top \mathbf{Q}_{ii} \tilde{\mathbf{x}}_i - 2 \sum_{i=1}^n \sum_{j \neq i} \tilde{\mathbf{x}}_i^\top \mathbf{Q}_{ij} \tilde{\mathbf{x}}_j \\ & + 2 \sum_{i=1}^n \sum_{j \neq i} \tilde{\mathbf{x}}_i^\top \mathbf{Q}_{ij} (\mathbf{e}_{(j,i)}) \\ & - 2 \sum_{i=1}^n \sum_{j \neq i} (\mathbf{e}_{(i,j)})^\top \left(\sum_{k=1}^n \mathbf{Q}_{ik} \tilde{\mathbf{x}}_k \right) \\ & + 2 \sum_{i=1}^n \sum_{j \neq i} \sum_{k \neq i} (\mathbf{e}_{(i,j)})^\top \mathbf{Q}_{ik} (\mathbf{e}_{(k,i)}) \\ & - 2 \sum_{i=1}^n \sum_{j \neq i} (\mathbf{e}_{(j,i)})^\top a_{(j,i)}(t) \tilde{\mathbf{x}}_j \\ & + \sum_{i=1}^n \sum_{j \neq i} (2a_{(j,i)}(t) - \lambda_{(j,i)}) (\mathbf{e}_{(j,i)})^\top (\mathbf{e}_{(j,i)}) \\ & - 2 \sum_{i=1}^n \sum_{j \neq i} (\mathbf{e}_{(j,i)})^\top a_{(j,i)}(t) \mathbf{y}_j^* \\ & \stackrel{(e)}{=} W_1 + 2W_2 + W_3 + 2W_4 + 2W_5 + W_6 + W_7 \quad (15) \end{aligned}$$

with expressions

$$W_1 = -2\tilde{\mathbf{x}}^\top \mathbf{Q} \tilde{\mathbf{x}} \quad (17a)$$

$$W_2 = \sum_{i=1}^n \sum_{j \neq i} (\mathbf{e}_{(j,i)})^\top \mathbf{Q}_{ji} \tilde{\mathbf{x}}_i \quad (17b)$$

$$W_3 = 2 \sum_{i=1}^n \sum_{j \neq i} \sum_{k \neq j} (\mathbf{e}_{(j,i)})^\top \mathbf{Q}_{jk} (\mathbf{e}_{(k,j)}) \quad (17c)$$

$$W_4 = - \sum_{i=1}^n \sum_{j \neq i} (\mathbf{e}_{(j,i)})^\top \left(\sum_{k=1}^n \mathbf{Q}_{jk} \tilde{\mathbf{x}}_k \right) \quad (17d)$$

$$W_5 = - \sum_{i=1}^n \sum_{j \neq i} (\mathbf{e}_{(j,i)})^\top a_{(j,i)}(t) \mathbf{I}_{d_j} \tilde{\mathbf{x}}_j \quad (17e)$$

$$W_6 = \sum_{i=1}^n \sum_{j \neq i} (\mathbf{e}_{(j,i)})^\top (2a_{(j,i)}(t) - \lambda_{(j,i)}) \mathbf{I}_{d_j} (\mathbf{e}_{(j,i)}) \quad (17f)$$

$$W_7 = -2 \sum_{i=1}^n \sum_{j \neq i} (\mathbf{e}_{(j,i)})^\top a_{(j,i)}(t) \mathbf{I}_{d_j} \mathbf{y}_j^* \quad (17g)$$

using relabelled counters where we,

- (a) apply the Itô jump rule, carefully considering the partial derivatives of the Lyapunov function and the value of V immediately following a jump event and insert (12),
- (b) substitute in the error dynamics derived from Lemma 2,
- (c) expand the first term of (b) and introduce zero-sum terms to enhance clarity,
- (d) expand and combine expressions, explicitly incorporating the newly defined error symbols, such as $\mathbf{e}_{(i,p)}$, for improved readability and conciseness and
- (e) group the resulting bilinear and quadratic forms, thus expressing the total infinitesimal generator as a summation over individual contributions.

Additionally, we upper-bound the bilinear expression W_7 with the intention to absorb it partly into the bilinear terms as

$$W_7 \stackrel{(a)}{=} -2 \sum_{i=1}^n \sum_{j \neq i} (a_{(j,i)}(t) \mathbf{I}_{d_j} \mathbf{e}_{(j,i)})^\top \mathbf{y}_j^*$$

$$\begin{aligned}
&\stackrel{(b)}{\leq} \sum_{i=1}^n \sum_{j \neq i} \left(\rho_j \|a_{(j,i)}(t) \mathbf{I}_{d_j} \mathbf{e}_{(j,i)}\|^2 + \frac{1}{\rho_j} \|\mathbf{y}_j^*\|^2 \right) \\
&\stackrel{(c)}{=} \sum_{i=1}^n \sum_{j \neq i} \left(\rho_j a_{(j,i)}(t)^2 \|\mathbf{e}_{(j,i)}\|^2 + \frac{1}{\rho_j} \|\mathbf{y}_j^*\|^2 \right). \quad (17h)
\end{aligned}$$

where we

- (a) plug in the original expression for W_7 from (17g), with the drift rates $a_{(j,i)}$ grouped with the error terms $\mathbf{e}$,
- (b) apply Young's inequality introduce positive weighting constants $(\rho_j)_{j=1,\dots,n}$ and
- (c) expand the norm using submultiplicativity of norms.

We are now ready to reformulate an upper bound for the infinitesimal generator as a bilinear form.

(Step 3) The infinitesimal generator can be upper-bounded by the quadratic form $-\mathbf{s}^\top \mathbf{M}(t) \mathbf{s} + \gamma'$. The matrix $\mathbf{M}(t)$, encapsulating the quadratic and bilinear terms, takes the block structure

$$\mathbf{M}(t) = \begin{bmatrix} \mathbf{M}_{11} & \mathbf{M}_{21}^\top(t) \\ \mathbf{M}_{21}(t) & \mathbf{M}_{22}(t) \end{bmatrix}.$$

The blocks of $\mathbf{M}(t)$ are formed by systematically collecting the coefficients from the expanded expression of $\mathcal{LV}$ from (16), incorporating the upper bound for W_7 from (17h).

The block $\mathbf{M}_{11}$, capturing quadratic terms involving only $\tilde{\mathbf{x}}$, is given by

$$\mathbf{M}_{11} = 2\mathbf{Q} \quad (18)$$

The block $\mathbf{M}_{21}(t)$ captures bilinear terms involving products of $\mathbf{e}$ and $\tilde{\mathbf{x}}$, defined by contributions from W_2 , W_4 , and W_5 as

$$\begin{aligned}
(\mathbf{M}_{21})_{(j,i),(k)}(t) &= \mathbf{Q}_{jk} - \begin{cases} \mathbf{Q}_{ji} & \text{if } k = i \\ \mathbf{0} \in \mathbb{R}^{d_i \times d_k} & \text{if } k \neq i \end{cases} \\
&\quad + \begin{cases} a_{(j,i)}(t) \mathbf{I}_{d_j} & \text{if } k = j \\ \mathbf{0} \in \mathbb{R}^{d_i \times d_k} & \text{if } k \neq j \end{cases} \quad (19)
\end{aligned}$$

The block $\mathbf{M}_{22}(t)$ comprises quadratic terms solely in $\mathbf{e}$. As per the statement of the theorem, we express $\mathbf{M}_{22}(t)$ as the sum of a block-diagonal matrix $\mathbf{\Lambda}$ containing the Poisson rates and a matrix $\mathbf{R}(t)$ containing remaining terms. Thus, $\mathbf{M}_{22}(t) = \mathbf{\Lambda} + \mathbf{R}(t)$, where

$$\mathbf{\Lambda} = \text{diag}(\dots, \lambda_{(j,i)} \mathbf{I}_{d_j}, \dots), \quad (20a)$$

$$(\mathbf{R})_{((j,i),(p,q))}(t) = \quad (20b)$$

$$\begin{cases} (-2a_{(j,i)}(t) - \rho_j a_{(j,i)}(t)^2) \mathbf{I}_{d_j} & \text{if } (j,i) = (p,q) \\ -2\mathbf{Q}_{jp} & \text{if } i = q \text{ and } j \neq p \\ \mathbf{0} & \text{otherwise} \end{cases}$$

The diagonal terms of $\mathbf{R}(t)$ incorporate drift rates $a_{(j,i)}$ and weighting parameters ρ_j , while off-diagonal terms reflect cross-coupling from $\mathbf{Q}$.

The constant γ' aggregates all terms independent of $\mathbf{s}$, arising solely from the upper bound of W_7 , where

$$\gamma' = \sum_{i=1}^n \sum_{j \neq i} \frac{1}{\rho_j} \|\mathbf{y}_j^*\|^2 \quad (21)$$

We must show that for any given $\gamma > 0$ and $c_3 > 0$ (e.g., $c_3 = \beta$), constants $\rho_j > 0$ exist such that $\gamma' \leq c_3 \gamma$. This condition is $\sum_{i=1}^n \sum_{j \neq i} \frac{1}{\rho_j} \|\mathbf{y}_j^*\|^2 \leq c_3 \gamma$. As long as $\mathbf{y}^* \neq \mathbf{0}$ (if $\mathbf{y}^* = \mathbf{0}$, $\gamma' = 0$ and Corollary 1 applies), we can always find such ρ_j . For example, choosing $\rho_j = \rho$ for all j , the condition

becomes $\rho \geq \frac{1}{c_3 \gamma} (n-1) \|\mathbf{y}^*\|^2$. Since the right-hand side is a finite positive constant, such ρ_j always exist. We now proceed assuming such ρ_j have been chosen. This fixes γ' and also defines the matrix $\mathbf{R}(t)$ which depends on ρ_j .

(Step 4) At any fixed instant t , $\mathbf{M}(t)$ is a symmetric matrix. We apply Lemma 3 to this matrix for both statements of this theorem.

- (a) We need to show $\mathbf{M}(t) \succ \mathbf{0}$. By Lemma 3, this holds if all communication rates satisfy

$$\lambda_{(j,i)} > \lambda_s(t) := -\mu_{\min}(\mathbf{R}(t) - \mathbf{K}(t)), \quad (22a)$$

where $\mathbf{K}(t) := \mathbf{M}_{21}^\top(t) \mathbf{M}_{11}^{-1} \mathbf{M}_{21}(t)$, and consequently, $\mathcal{LV}(\mathbf{s}) \leq -\mu_{\min}(\mathbf{M}(t)) \|\mathbf{s}\|^2 + \gamma'$. We set $c_3(t) = \mu_{\min}(\mathbf{M}(t)) > 0$.

- (b) We need to show that $\mathbf{M}(t) \succeq \beta \mathbf{I}$ for the given $\beta \in (0, 2\mu_{\min}(\mathbf{Q}))$. Let $\mu = \beta$. Since $\mathbf{M}_{11} = 2\mathbf{Q}$, $\mu_{\min}(\mathbf{M}_{11}) = 2\mu_{\min}(\mathbf{Q})$ and the condition $\mu \in (0, \mu_{\min}(\mathbf{M}_{11}))$ from Lemma 3 is satisfied. By Lemma 3, $\mathbf{M}(t) \succeq \beta \mathbf{I}$ holds, if all $\lambda_{(j,i)}$ satisfy

$$\lambda_{(j,i)} \geq \lambda_d(t) := \beta - \mu_{\min}(\mathbf{R}(t) - \mathbf{K}_\beta(t)), \quad (22b)$$

where $\mathbf{K}_\beta(t) := \mathbf{M}_{21}^\top(t) (\mathbf{M}_{11} - \beta \mathbf{I})^{-1} \mathbf{M}_{21}(t)$, and consequently $\mathcal{LV}(\mathbf{s}) \leq -\beta \|\mathbf{s}\|^2 + \gamma'$. We set $c_3 = \beta$.

To satisfy Lemma 1 and guarantee stability for all time, we require a constant lower bound $c_3 \in (0, \infty)$ such that $\mathcal{LV}(\mathbf{s}) \leq -c_3 \|\mathbf{s}\|^2 + \gamma'$. This requires finding constant rates λ_s and λ_d that satisfy the worst-case conditions $\lambda_{(j,i)} > \lambda_s \geq \sup_t \{\lambda_s(t)\}$ (for (a)) and $\lambda_{(j,i)} \geq \lambda_d \geq \sup_t \{\lambda_d(t)\}$ (for (b)).

Finding the exact infimum of $\mu_{\min}(\cdot)$ over time is intractable. Instead, we construct constant matrices $\mathbf{R}_c$, $\mathbf{K}_c$, and $\mathbf{K}_{c,\beta}$ (hereafter $\mathbf{R}$, $\mathbf{K}$, $\mathbf{K}_\beta$) that serve as uniform bounds, such that $\mathbf{R}(t) \geq \mathbf{R}$, $\mathbf{K}(t) \leq \mathbf{K}$, and $\mathbf{K}_\beta(t) \leq \mathbf{K}_\beta$ for all t .

This allows us to bound the infima, as $\inf_t \{\mu_{\min}(\mathbf{R}(t) - \mathbf{K}(t))\} \geq \mu_{\min}(\mathbf{R} - \mathbf{K})$ and $\inf_t \{\mu_{\min}(\mathbf{R}(t) - \mathbf{K}_\beta(t))\} \geq \mu_{\min}(\mathbf{R} - \mathbf{K}_\beta)$.

· **Bounding $\mathbf{R}(t)$:** We require a constant matrix $\mathbf{R}$ such that $\mathbf{R}(t) \geq \mathbf{R}$. The diagonal entries of $\mathbf{R}(t)$ are $-2a_{(j,i)}(t) - \rho_j a_{(j,i)}(t)^2$. This is a downward-opening parabola in $a_{(j,i)}(t)$. Given the assumption $|a_{(j,i)}(t)| \leq a_{ji}$, its minimum on the interval $[a_{ji}, a_{ji}]$ occurs at a boundary. Thus, we define $\mathbf{R}$ as the constant matrix resulting from replacing $a_{(j,i)}(t)$ with a_{ji} in (20b).

· **Bounding $\mathbf{K}(t)$ and $\mathbf{K}_\beta(t)$:** We require constant matrices $\mathbf{K}$, $\mathbf{K}_\beta$ such that $\mathbf{K}(t) \leq \mathbf{K}$ and $\mathbf{K}_\beta(t) \leq \mathbf{K}_\beta$ for all t . We find a valid upper bound $\mathbf{K}_\beta$ by bounding its maximum eigenvalue $\mu_{\max}(\mathbf{K}_\beta(t)) \leq \|\mathbf{K}_\beta(t)\|_2$. Using the submultiplicativity of matrix norms, we have $\|\mathbf{K}_\beta(t)\|_2 \leq \|\mathbf{M}_{21}(t)\|_2^2 \cdot \|(\mathbf{M}_{11} - \beta \mathbf{I})^{-1}\|_2 = \|\mathbf{M}_{21}(t)\|_2^2 \cdot \|(2\mathbf{Q} - \beta \mathbf{I})^{-1}\|_2$. Next, we decompose $\mathbf{M}_{21}(t) = \mathbf{M}_{21,c} + \mathbf{M}_{21,v}(t)$, where $\mathbf{M}_{21,c}$ contains the constant $\mathbf{Q}$ -terms and $\mathbf{M}_{21,v}(t)$ contains the time-dependent $a_{(j,i)}(t)$ terms. We apply the triangle inequality $\|\mathbf{M}_{21}(t)\|_2 \leq \|\mathbf{M}_{21,c}\|_2 + \|\mathbf{M}_{21,v}(t)\|_2$. The norm of the constant part $\|\mathbf{M}_{21,c}\|_2$ is a fixed value. The norm of the (block-)diagonal variable part $\mathbf{M}_{21,v}(t)$ is bounded by its largest entry in magnitude, using $|a_{(j,i)}(t)| \leq a_{ji}$ as $\|\mathbf{M}_{21,v}(t)\|_2 = \max_{(j,i)} \{|-a_{(j,i)}(t)|\} \leq \max_{(j,i)} \{a_{ji}\} =: a_{\max}$. Combining these steps, we obtain a constant upper bound $K_{\text{bound},\beta}$ for the maximum eigenvalue of $\mathbf{K}_\beta(t)$ as $K_{\text{bound},\beta} := (\|\mathbf{M}_{21,c}\|_2 + a_{\max})^2 \cdot \|(2\mathbf{Q} - \beta \mathbf{I})^{-1}\|_2$. We define $\mathbf{K}_\beta = K_{\text{bound},\beta} \mathbf{I}$. This constant matrix satisfies $\mathbf{K}_\beta(t) \leq \mu_{\max}(\mathbf{K}_\beta(t)) \mathbf{I} \leq \mathbf{K}_\beta$. The identical procedure, setting $\beta = 0$,

yields the bound $\mathbf{K} = K_{\text{bound}}\mathbf{I}$ for $\mathbf{K}(t)$, where K_{bound} uses $\|(2\mathbf{Q})^{-1}\|_2$.

With these constant bounds, we can define the sufficient constant rates

$$\lambda_s := -\mu_{\min}(\mathbf{R} - \mathbf{K}), \quad \text{and} \quad (23a)$$

$$\lambda_d := \beta - \mu_{\min}(\mathbf{R} - \mathbf{K}_\beta). \quad (23b)$$

(Step 5) Finally, we verify the parameters of Lemma 1. We chose $c_1 = c_2 = 1$. In Step 2, we showed that for any given $\gamma > 0$ and $c_3 > 0$ (where $c_3 = \mu_{\min}(\mathbf{M}_c)$ for (a) or $c_3 = \beta$ for (b)), we can choose $\rho_j > 0$ such that the resulting $\gamma' = (n-1) \|\mathbf{y}^*\|_{P^{-1}}^2$ satisfies $\gamma' \leq c_3\gamma$. The parameters from Lemma 1 are $\alpha = c_2/c_1 = 1$, $\beta = c_3/c_2 = c_3$ (which is β_{Lemma} for statement (b)), $\gamma_{\text{Lemma}} = \gamma'/c_3 \leq \gamma$. This confirms that the system is stable with the desired exponential rate (for (b)) and the desired ultimate bound γ .

□

Proof (of Corollary 1) Observe $W_7 \equiv 0$ if $a_{p \rightarrow i} = 0$ for all channels (p, i) . Consequently, $\gamma' = 0$ in the proof of Theorem 1.

□