
Distributed Event-Based Learning via ADMM

Guener Dilsad ER¹ Sebastian Trimpe² Michael Muehlebach¹

Abstract

We consider a distributed learning problem, where agents minimize a global objective function by exchanging information over a network. Our approach has two distinct features: (i) It substantially reduces communication by triggering communication only when necessary, and (ii) it is agnostic to the data-distribution among the different agents. We therefore guarantee convergence even if the local data-distributions of the agents are arbitrarily distinct. We analyze the convergence rate of the algorithm both in convex and nonconvex settings and derive accelerated convergence rates for the convex case. We also characterize the effect of communication failures and demonstrate that our algorithm is robust to these. The article concludes by presenting numerical results from distributed learning tasks on the MNIST and CIFAR-10 datasets. The experiments underline communication savings of 35% or more due to the event-based communication strategy, show resilience towards heterogeneous data-distributions, and highlight that our approach outperforms common baselines such as FedAvg, FedProx, SCAF-FOLD and FedADMM.

1. Introduction

Distributed learning refers to the minimization of a global objective function over a network of agents, where each agent has only access to a local cost function and can communicate with some or all agents in the network. Distributed learning systems provide a solution for handling the growing amount of data being generated everywhere on earth, by utilizing the computational power of individual devices in a network rather than relying on a central entity. This takes

¹Learning and Dynamical Systems Group, Max Planck Institute for Intelligent Systems, 72076 Tübingen, Germany ²Institute for Data Science in Mechanical Engineering, RWTH Aachen University, 52068 Aachen, Germany. Correspondence to: Guener Dilsad ER <gder@tue.mpg.de>.

the burden off central processors and improves data privacy by avoiding a centralized training and storage of data.

Distributed learning is particularly challenging when the data is not independent and identically distributed (non-i.i.d.) across the different agents. This situation often hinders the convergence to a globally optimal model. The non-i.i.d. nature leads to disparities in local datasets, preventing the local models from generalizing across the entire dataset, leading to a fundamental dilemma between minimizing local and global objective functions (Acar et al., 2021). In addition, a second key challenge arises from the communication between agents, which is required to ensure convergence to the global solution and may lead to a substantial overhead. This communication overhead results in a waste of energy (Li et al., 2020b), and is prone to delays and communication channel failures. As a result, both, non-i.i.d. datasets and communication overhead, constitute major bottlenecks for enabling large-scale learning systems.

We provide an effective solution to both challenges. Inspired by the sent-on-delta concept (Miskowicz, 2006), we reduce the communication load by introducing an event-based communication strategy, such that each agent (or computational node) communicates only if necessary. Our communication rule enforces local constraints that collectively guarantee bounded overall error, a paradigm related to safe zone design strategies in distributed computing (Garofalakis & Samoladas, 2017). Our approach is also rooted in event-based estimation, where communication is triggered by significant state changes. We further base our approach on the Alternating Direction Method of Multipliers (ADMM). Our method is therefore robust against ill-conditioning and agnostic towards a disparity of the local data-distributions among the agents (these can be skewed in arbitrary ways). The approach further enables an explicit trade-off between communication load on the network and solution accuracy via a small set of hyperparameters that have a clear interpretation. We explicitly quantify the influence of these hyperparameters on the solution accuracy and analyze the effect of communication failures. The article concludes by highlighting the effectiveness of our algorithm in training neural networks, and solving LASSO problems in a distributed and communication-efficient manner.

Our theoretical analysis builds on a recent trend in the opti-

mization literature (Wibisono et al., 2016; Su et al., 2016; Muehlebach & Jordan, 2019; Tong & Muehlebach, 2023) that views algorithms as dynamical systems and leverages ideas from differential or symplectic geometry, as well as passivity and dissipativity (Lessard et al., 2016; Muehlebach & Jordan, 2020). As we will show, this enables convergence proofs and a convergence rate analysis for our distributed algorithms, together with an analysis of robustness against communication failures. Our work provides important insights into the behavior of the event-based optimization under communication failures, an aspect, which has been overlooked in prior works, and thereby lays the groundwork for future research in this area.

Related Work: In the 1980s, Bertsekas & Tsitsiklis (1989) and others laid the foundation for the analysis of distributed algorithms. As machine learning became popular, distributed learning emerged, specifically focusing on parallelizing computation for empirical risk minimization. Shokri & Shmatikov (2015) explored collaborative deep learning with multiple agents using distributed stochastic gradient descent, which was later coined federated learning (McMahan et al., 2017) and advanced by subsequent contributions (Kairouz et al., 2021; Asad et al., 2023). A unifying element in these works is the consensus problem, where agents agree on a common value or decision. This problem is both central to distributed optimization and is also a special instance of distributed optimization (Wei & Ozdaglar, 2012).

The trade-off between communication and computation is inevitable in distributed optimization (Nedic et al., 2018). Recent work by Cao et al. (2023) categorizes communication-efficient distributed learning into four main strategies: (1) minimizing the number of communications, (2) compression, (3) managing resources (e.g., bandwidth), and (4) using game theoretical approaches. We focus our review on the first category that aligns with our work, and reduces communication by transmitting information only if necessary. A first line of work (McMahan et al., 2017; Wei Liu et al., 2021; Reisizadeh et al., 2020, and many more) proposes algorithms with a periodical exchange of model parameters either among all agents or randomly selected subsets for decreasing communication load. While this approach is particularly straightforward and easy to implement, the random sampling risks missing critical updates or performing redundant communications. A second line of work involves accelerated gradient methods for distributed optimization, reducing the need for many communication rounds to converge. For instance, Kovalev et al. (2020) and Nabli & Oyallon (2023) optimize the number of gradient evaluations together with communication rounds. Shamir et al. (2014) replaces gradient descent with Newton-like methods and Hendrikx et al. (2020) proposes statistical preconditioning where both methods further improve convergence rates at the cost of a higher computational load per iteration. Addi-

tionally, Liu et al. (2021) propose a lazy evaluation of dual gradients, reducing communication by skipping redundant updates, while (Chen et al., 2018) adaptively reuse lagged gradients to meet target accuracy with fewer communication rounds. There has also been a third line of work that focuses on reducing communication via event-based triggering and compression of network parameters (Liu et al., 2019; Ghadikolaei et al., 2021; Singh et al., 2023; Zhang et al., 2023). While (Zhang et al., 2024; 2023) employ an ADMM-based strategy that is similar to ours, their focus lies on investigating different compression schemes, and not on analyzing convergence rates and the effect of communication failures. Event-triggering has also been explored in contexts like dynamics model learning (Solowjow & Trimpe, 2020; Umlauf & Hirche, 2019), and Bayesian optimization (Brunzema et al., 2025). While highlighting the benefit of triggering for reducing communication, these works do not consider distributed optimization problems as we do herein.

In addition to the communication overhead, another major challenge for distributed learning arises from non-i.i.d. data distributions across agents (Zhao et al., 2018; Li et al., 2020c; Glasgow et al., 2022). SCAFFOLD (Karimireddy et al., 2020) addresses this challenge by introducing a client control variate to improve convergence at the cost of doubling communication. Similarly, (Gao et al., 2022) enhances training with auxiliary drift variables, while (Zheng et al., 2024) selects representative clients and adjusts server gradients. Recent contributions by Li et al. (2020a); Acar et al. (2021); Shi et al. (2023) add a proximal regularization term to the local objective functions of the individual agents, whereas Zhang et al. (2021) (FedPD) and Zhou & Li (2023); Wang et al. (2022); Gong et al. (2022) (FedADMM) address the challenge with ADMM formulations. However, compared to our work, FedADMM (Zhou & Li, 2023; Wang et al., 2022; Gong et al., 2022) relies on utilizing a random selection of agents that communicate and FedPD (Zhang et al., 2021) considers full participation, whereas we use an event-triggered mechanism. Alternatively, other splitting schemes such as Douglas-Rachford method proposed by Tran Dinh et al. (2021), similarly align local and global objectives but remain constrained by random agent participation.

ADMM remains a widely-used tool for distributed learning, with recent advancements focusing on improving convergence rates and communication efficiency. For example, Wang et al. (2025) introduce inertia and adaptive iteration strategies to accelerate convergence, while Song et al. (2025) controls inexactness and dynamically tunes penalty parameters. In addition, He et al. (2023) explore dynamic tuning of ADMM hyperparameters, and hierarchical grouping approaches (Qiu et al., 2023), inspired by Elgabri et al. (2020), aim to reduce communication overhead by restricting updates to neighboring workers. These methods share the goal

of improving the efficiency of ADMM in distributed settings, but they still rely on periodic or full-agent participation, whereas our approach uses event-triggered mechanisms to further reduce communication costs.

As we also highlight in numerical experiments, a random selection of agents might prevent important local changes from propagating quickly through the network, leading to a slower convergence. To the best of our knowledge, this is the first work to provide a convergence analysis of distributed learning with event-triggered communication that addresses key aspects, such as packet drop and the presence of non-i.i.d. data.

Contributions are summarized as follows:

(i) We propose an event-based communication scheme for distributed optimization, where a communication event is only triggered, when the current state has deviated by a predefined threshold Δ , indicating a significant change in the local decision variables. Therefore, our approach is effective in reducing communication overhead and can adapt to the limited communication resources in heterogeneous networks. Our method is also compatible with and complementary to gradient compression/quantization (Hegazy et al., 2024; Mao et al., 2022; Wang et al., 2018) and fair aggregation techniques (Zhu & Ling, 2021).

(ii) We characterize the effect of the communication threshold Δ on the solution accuracy and therefore quantify the trade-off between communication and solution accuracy. Compared to other ADMM-based approaches, such as (Zhang et al., 2021; Zhou & Li, 2023), our method is versatile, both in the selection of variables that are being communicated (which is important for reducing communication in practice), as well as the different problem formulations that we can address. In particular, our approach goes beyond the scope of consensus problems, and can solve generic constrained optimization problems, sparse regression and LASSO problems, perform robust principal component analysis (Candès et al., 2011), and solve distributed learning instances where the features but not the data points are distributed (Boyd et al., 2010).

(iii) Numerical experiments support the theoretical analysis and highlight that our approach even converges in the most extreme non-i.i.d. setting, where each agent has only access to training data from a single class (see the MNIST classifier example in Sec. 5). Comparisons to the baselines FedADMM (Zhou & Li, 2023), SCAFFOLD (Karimireddy et al., 2020), FedProx (Li et al., 2020a) and FedAvg (McMahan et al., 2017) demonstrate superiority both in terms of communication efficiency and classification accuracy.

(iv) We demonstrate an accelerated convergence rate, and derive symbolic expressions that relate the convergence rate to instance-specific quantities such as the condition number

and the topology of the communication network. The convergence analysis requires a Lyapunov-like function that is different compared to earlier work (Nishihara et al., 2015), due to the presence of the event-based communication.

(v) We study the robustness of our algorithm against communication failures, both in theory as well as in numerical experiments, which, to the best of our knowledge, is largely missing in the literature (a notable exception for the consensus problem is (Bastianello et al., 2021)). We address communication failures algorithmically by proposing a rare periodic reset strategy. We show that, without such a reset strategy, inter-agent errors accumulate rapidly in the presence of packet drops and prevent convergence.

Outline: The article is structured as follows: Sec. 2 describes the problem formulation and introduces our event-based learning algorithm in the consensus setting. The more general formulation is discussed in Sec. 3, where we also introduce a dynamical systems model for our algorithm. Sec. 4 discusses the convergence analysis of the proposed algorithm and presents convergence rates, while empirical results that underline the theoretical findings are included in Sec. 5 and in App. G. The appendix contains additional technical details about the communication structure in App. A and the details of the convergence analysis in App. C and D.

2. Event-Based Distributed Learning

We consider a distributed learning problem of the type $\min_{x \in \mathbb{R}^n} \sum_{i=1}^N f^i(x)$, where the overall cost function $f(x)$ is the sum of N individual, potentially nonsmooth functions. The different f^i typically arise from different training datasets stored on different computational nodes. In the most basic instance, our algorithm arises from the consensus formulation

$$\begin{aligned} \min_{x^1, \dots, x^N \in \mathbb{R}^n} \quad & \sum_{i=1}^N f^i(x^i) + g(z), \\ \text{subject to} \quad & x^i = z, \quad i = 1, \dots, N, \end{aligned} \quad (1)$$

where we impose the constraints $x^i = z$ by corresponding dual variables u^i . Thus, by guaranteeing constraint satisfaction, we can ensure consensus between the agents despite different local problems and, in particular, arbitrary non-i.i.d. data distributions among the computational nodes. In addition, we assume the function $f^i : \mathbb{R}^n \rightarrow \mathbb{R}$ to be smooth, while $g : \mathbb{R}^n \rightarrow \bar{\mathbb{R}}$ is allowed to be nonsmooth (g typically represents a regularizer) and maps to the extended real numbers.

Our event-based algorithm, stated in Alg. 1, works as follows. Each agent (or computational node) $i, i = 1, \dots, N$, has access to the local objective function f^i , its local solution x^i , its local multiplier u^i , and an estimate \hat{z}^i of the consensus variable z . We further introduce the agent $N+1$

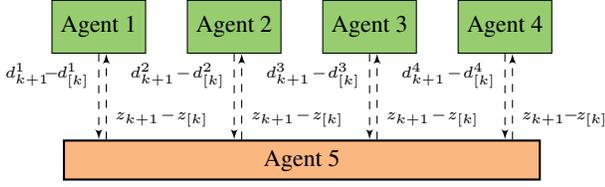


Figure 1: The figure illustrates the distributed learning setup. The Agents 1–4 store x^i, u^i and perform updates based on the information received by Agent 5, according to Alg. 1. Agent 5, stores z and performs updates based on the information received by Agent 1–4. This architecture is common in distributed learning, where a single server aggregates updates from multiple distributed clients to collaboratively train a model.

(acting as server) that has access to g , the variable z , and maintains an estimate $\hat{\zeta}$ of the average

$$\hat{\zeta}_k := \frac{1}{N} \sum_{i=1}^N (\alpha x_{k+1}^i + u_k^i).$$

Following the communication structure in Fig. 1, the algorithm proceeds in two steps:

i) Parallel update of agents $i = 1, \dots, N$: Each agent $i = 1, \dots, N$ first updates its estimate \hat{z}_k^i based on whether it receives an event-based communication from the agent $N+1$. The agent then updates its multiplier u_k^i and solves a local minimization over f^i , which also includes a quadratic regularization term. The regularization term ensures that the minimization is well-conditioned (a key advantage to dual ascent, for example) and the local solution x^i is biased towards \hat{z}_k^i . In practice, the minimization is replaced by a fixed number of (stochastic) gradient descent steps. If the resulting value $d_{k+1}^i := \alpha x_{k+1}^i + u_k^i$ of the agent i is significantly different from the value that it last communicated to the agent $N+1$, an event-based communication is triggered and the difference of d_{k+1}^i to the last communicated value is sent to the agent $N+1$.

ii) Update of agent $N+1$: The agent $N+1$ updates its estimate $\hat{\zeta}$ of ζ by accumulating the d^i variables that it receives from all agents. It then updates the consensus variable z_{k+1} by solving a local minimization over g with a quadratic regularization term. Note that if the nonsmooth component g is missing, z_{k+1} is simply set to $\hat{\zeta}_k - (1-\alpha)z_k$. Finally, the agent $N+1$ triggers an event-based communication if the value z_{k+1} is significantly different from the value that it last communicated to the agents $i = 1, \dots, N$.

Next, we explain the details of the event-based communication protocol on the example of the communication of d^i , which is related to the primal x^i and dual u^i variables of agent i . The other event-based communications proceed similarly. The protocol comes in two variants, **vanilla event-**

Algorithm 1 Event-Based Distributed Learning with Over-Relaxed ADMM

Require: Local objective functions f^i , parameters ρ, Δ^d, Δ^z , reset period T

Require: Initialize $\hat{x}_0^i = x_0, \hat{z}_0 = \zeta_0 = x_0, \hat{u}_{-1}^i = u_0^i$

for $k = 0$ to t_{\max} **do**

for $i = 1$ to N **do**

$\hat{z}_k^i \leftarrow$ receive $z_k - z_{[k-1]}$ {Agent i }
 $u_k^i = u_{k-1}^i + \alpha x_k^i - \hat{z}_k^i + (1-\alpha)\hat{z}_{k-1}^i$
 $x_{k+1}^i = \arg \min_{x^i} f^i(x^i) + \frac{\rho}{2} |x^i - \hat{z}_k^i + u_k^i|^2$
 event-based send of $d_{k+1}^i - d_{[k]}^i$ {See (2)}

end for

$\hat{\zeta}_k \leftarrow$ receive $\frac{1}{N} \sum_{i \in \mathcal{C}_{k+1}^d} (d_{k+1}^i - d_{[k]}^i)$ {Agent $N+1$ }

$z_{k+1} = \arg \min_z g(z) + \frac{N\rho}{2} |z - \hat{\zeta}_k - (1-\alpha)z_k|^2$
 event-based send of $z_{k+1} - z_{[k]}$

if $\text{mod}(k+1, T) = 0$ **then**

 perform reset $\rightarrow \hat{\zeta}_k = \zeta_k, \hat{z}_k = z_k$

end if

end for

based and randomized event-based.

Vanilla event-based: This communication rule is inspired from the sent-on-delta concept (Miskowicz, 2006), which aims to reduce the number of communications by only sending updates when significant changes occur. A communication is triggered, if the value d_{k+1}^i has deviated by more than the predefined threshold $\Delta^d > 0$ compared to the value that was last communicated. We introduce the variable $d_{[k]}^i$ to denote the value d_k^i that was last communicated and add the index i to the set \mathcal{C}_{k+1}^d . The set \mathcal{C}_{k+1}^d denotes the set of agents that trigger a communication of d_{k+1}^i at time-step k , that is,

$$|d_{k+1}^i - d_{[k]}^i| > \Delta^d \iff i \in \mathcal{C}_{k+1}^d, \quad (2)$$

and $d_{k+1}^i - d_{[k]}^i$ is sent out. Similarly, agent $N+1$ triggers a communication if $|z_{k+1} - z_{[k]}| > \Delta^z$. We also model communication failures as drops, which we represent by the variables χ_{k+1}^{di} . The variable χ_{k+1}^{di} takes the value $\chi_{k+1}^{di} = -(d_{k+1}^i - d_{[k]}^i)$, if $d_{k+1}^i - d_{[k]}^i$ is not received by the agent $N+1$; otherwise $\chi_{k+1}^{di} = 0$. The agent $N+1$ updates its estimate of the average $\hat{\zeta}_k$ according to the primal and dual variables that it has received at time k , that is,

$$\hat{\zeta}_k = \hat{\zeta}_{k-1} + \frac{1}{N} \sum_{i \in \mathcal{C}_{k+1}^d} (d_{k+1}^i - d_{[k]}^i + \chi_{k+1}^{di}).$$

Randomized event-based: The protocol makes a case distinction. If $|d_{k+1}^i - d_{[k]}^i| \leq \Delta^d$, a communication is randomly triggered with probability p_{trig} . If $|d_{k+1}^i - d_{[k]}^i| > \Delta^d$, a communication is triggered with certainty. Randomized communication from agent $N+1$ to the other agents works in a similar way. If $|z_{k+1} - z_{[k]}| \leq \Delta^z$, then a

communication is randomly triggered with probability p_{trig} between agent $N + 1$ and agent i .

We observed in our numerical experiments that **randomized event-based** often improves **vanilla event-based** in terms of the achieved communication versus solution accuracy trade-off.

The error caused by the event-based communication remains bounded at all times thanks to the communication protocol and the periodic resets. This is summarized with the next proposition, whose proof is included in App. E:

Proposition 2.1. *The error $\hat{\zeta}_k - \zeta_k$ at iteration k is bounded by $|\hat{\zeta}_k - \zeta_k| \leq \Delta^d + T\bar{\chi}^d$, where T denotes the reset period (see Alg. 1) and $\bar{\chi}^d$ is a bound on the disturbance χ_k^{di} .*

We now state the main convergence result for Alg. 1. The result arises as a corollary from the convergence analysis in the more general distributed optimization setting (see Thm. 4.1). We also provide sublinear convergence rates in a nonconvex setting (see Thm 2.3).

Corollary 2.2. *Let $f = \sum_{i=1}^N f^i$ be m -strongly convex and L -smooth with $\kappa = L/m$, and g be convex. Let the step-size be $\rho = (mL)^{\frac{1}{2}} \kappa^\epsilon$ with $\epsilon \in [0, \infty)$, and $\alpha = 1$. For large enough κ , we have*

$$|z_k - z_*|^2 \leq 4 \left(1 - \frac{1}{4\kappa^{\epsilon + \frac{1}{2}}}\right)^{2k} D_0 + \frac{5}{N} \kappa^{2+2\epsilon} \Delta^2,$$

where z_* is the optimal value for the consensus variable z , and D_0 represents the initial error; $D_0 = |z_0 - z_*|^2 + \frac{1}{N} \sum_{i=1}^N |u_0^i - u_*^i|^2$, with u_*^i denoting the optimal values of the dual variables associated with each agent. Here, $\Delta = N\Delta^d + \Delta^z + T(N\bar{\chi}^d + \bar{\chi}^z)$ captures the error arising from the event-based communication.

The convergence result bounds the distance between the consensus variable z_k and the optimal solution $z_* = x_*$ that minimizes (1). The analysis models our event-based learning algorithm as a dynamical system, accounting for disturbances introduced by the event-based communication strategy. By design, these disturbances remain bounded under the communication protocol. The next section elaborates on the formulation of our algorithm as a dynamical system.

The strong convexity assumption enables faster convergence rates compared to more general nonconvex scenarios. Specifically, under this assumption, the rate of convergence is linear, as shown in Cor. 2.2 and accelerated. In contrast, without such assumptions, convergence rates are generally much slower, typically sublinear or achieving only asymptotic convergence.

We note that the strong convexity assumption in Cor. 2.2 only requires $f := \sum_{i=1}^N f^i$ to be strongly convex, without

imposing the same condition on the individual components f^i . In addition, we present a convergence result for general nonconvex cases in Thm. 2.3 leading to sublinear convergence rates. The proof is provided in App. B.

Theorem 2.3. *Let each $f^i : \mathbb{R}^n \rightarrow \mathbb{R}$ be smooth (potentially nonconvex) and let $g : \mathbb{R}^n \rightarrow \bar{\mathbb{R}}$ be a proper, closed convex function. Let the relaxation parameter be $\alpha = 1$, and the communication threshold Δ_k decay as $\Delta_k = \Delta_0/(k+1)^2$. Then, the gradients and residuals converge with a rate of $\mathcal{O}(1/k)$, and the following bound holds:*

$$\frac{1}{K+1} \sum_{k=0}^K \left(\frac{2}{3N} \sum_{i=1}^N |r_{k+1}^i|^2 + \frac{1}{6N} |G_{k+1}|^2 \right) = \mathcal{O}\left(\frac{1}{K}\right),$$

where $r_{k+1}^i = x_{k+1}^i - z_{k+1}$ are the residuals, and the gradient terms are given by

$$G_{k+1} \in \frac{1}{\rho N} \left(\sum_{i=1}^N \nabla f^i(x_{k+1}^i) + \partial g(z_{k+1}) \right).$$

3. Event-Based ADMM as a Dynamical System

We introduce a more general problem formulation that encompasses the previous section as a special case in order to broaden the scope of our analysis. This leads to the following constrained minimization problem

$$\min_{x \in \mathbb{R}^p, z \in \mathbb{R}^q} f(x) + g(z), \quad \text{subject to } Ax + Bz = c, \quad (3)$$

where $x \in \mathbb{R}^p$ and $z \in \mathbb{R}^q$ are decision variables, $A \in \mathbb{R}^{r \times p}$, $B \in \mathbb{R}^{r \times q}$, and $c \in \mathbb{R}^r$ are corresponding matrices, and the objective function is decomposed into a smooth part $f : \mathbb{R}^p \rightarrow \mathbb{R}$ and a nonsmooth part $g : \mathbb{R}^q \rightarrow \bar{\mathbb{R}}$. We will provide an analysis under the following standard assumptions in distributed optimization.

Assumption 3.1. The matrix A is invertible and B is full column rank.

Assumption 3.2. The function f is m -strongly convex and L -smooth. The function g is convex.

It is important to emphasize that the assumption of strong convexity for f is introduced to derive linear rates within a dynamical systems framework. This assumption does not limit the practical applicability of the algorithm and a corresponding nonconvex result is included in Thm. 2.3. We also note that Assumption 3.2 allows for nonconvex f^i (see (1)) that $\sum f^i$ is strongly convex.

The formulation in (3) accommodates a variety of distributed optimization problems, including consensus, resource sharing, and distributed model fitting, see for example (Boyd et al., 2010). App. A further highlights how the general formulation can be tailored and simplified to accommodate specific applications, such as the sharing problem or finding a consensus on a graph.

Our event-based distributed learning method is summarized in Alg. 2. The algorithm is based on an over-relaxed version of ADMM, where an event-based communication structure between different agents is introduced. The over-relaxation brings the additional parameter α , which, as we will show, can be used to achieve faster convergence rates. The communication structure of the algorithm is shown in Fig. 2 and includes three agents that keep track of the individual quantities $r_k = Ax_k$, $s_k = Bz_k$, and the dual multiplier u_k . In the special case of the consensus problem, the updates of the primal variable x_k and dual variable u_k decompose further into local updates based on x_k^i and u_k^i , which results in the communication structure shown in Fig. 1.

Alg. 2 begins by initializing its variables, and over a series of iterations, agents alternate between sharing information and optimizing their local variables. Key steps include updating variables based on local objectives and residuals, and triggering communication events when individual residual changes exceed predefined thresholds. The algorithm leverages event-based communication to reduce the communication load, while still achieving convergence towards an optimal solution of (3), as we will show in the following section. The event-based communication proceeds as in Sec. 2, that is, the r -agent, for example, triggers a communication with the other agents if $|r_{k+1} - r_{[k]}| > \Delta^r$, at which point it sends the difference $r_{k+1} - r_{[k]}$ to the other agents. We again model communication failures by introducing the variables χ_{k+1}^{ru} if the communication is not received by the u -agent at time $k+1$. The notation is analogous for the remaining agents and communication lines (see also Fig. 2).

Algorithm 2 Event-Based Distributed Optimization with Over-Relaxed ADMM

Require: Functions f and g , matrices A and B , vector c , parameters ρ and α . Initial condition x_0, z_0

$$r_0 = \hat{r}_0^s = \hat{r}_0^u = Ax_0, s_0 = \hat{s}_0^r = \hat{s}_0^u = Bz_0, u_0 = \hat{u}_0^r = \hat{u}_0^s = 0$$

for $k = 0$ to t_{\max} **do**

$\hat{s}_k^r, \hat{u}_k^r \leftarrow$ receive $s_{k+1} - s_{[k]}, u_{k+1} - u_{[k]}$
 $x_{k+1} = \operatorname{argmin}_x f(x) + \frac{\rho}{2} \|Ax + \hat{s}_k^r - c + \hat{u}_k^r\|^2$
 event-based send $r_{k+1} - r_{[k]}$ where $r_{k+1} = Ax_{k+1}$

$\hat{r}_{k+1}^s, \hat{u}_k^s \leftarrow$ receive $r_{k+1} - r_{[k]}, u_{k+1} - u_{[k]}$
 $z_{k+1} = \operatorname{argmin}_z g(z) + \frac{\rho}{2} \|\alpha \hat{r}_{k+1}^s - (1-\alpha)Bz_k + Bz - \alpha c + \hat{u}_k^s\|^2$
 event-based send $s_{k+1} - s_{[k]}$, where $s_{k+1} = Bz_{k+1}$

$\hat{r}_{k+1}^u, \hat{s}_{k+1}^u \leftarrow$ receive $r_{k+1} - r_{[k]}, s_{k+1} - s_{[k]}$
 $u_{k+1} = u_k + \alpha \hat{r}_{k+1}^u - (1-\alpha)\hat{s}_k^u + \hat{s}_{k+1}^u - \alpha c$
 event-based send $u_{k+1} - u_{[k]}$

if $\operatorname{mod}(k+1, T) = 0$ **then**
 reset $\rightarrow \hat{r}_{k+1}^{u;s} = r_{k+1}, \hat{s}_{k+1}^{u;r} = s_{k+1}, u_{k+1}^{r;s} = u_{k+1}$
end if

end for

Alg. 2 has three update steps that occur sequentially,

whereby the first two involve optimization problems that can be replaced by their corresponding stationarity conditions. This yields the following implicit update equations:

$$\begin{aligned} 0 &= \nabla f(x_{k+1}) + \rho A^\top (Ax_{k+1} + \hat{s}_k^r - c + \hat{u}_k^r) \\ 0 &\in \partial g(z_{k+1}) + \rho B^\top (\alpha \hat{r}_{k+1}^s - (1-\alpha)Bz_k + Bz_{k+1} - \alpha c + \hat{u}_k^s) \\ u_{k+1} &= u_k + \alpha \hat{r}_{k+1}^u - (1-\alpha)\hat{s}_k^u + \hat{s}_{k+1}^u - \alpha c, \end{aligned}$$

which can be expressed by the dynamical system shown in Fig. 2. We note that the variable x_{k+1} is uniquely determined by \hat{s}_k^r and \hat{u}_k^r and does not depend on x_k , which means that only $\xi_k := (s_k, u_k)$ comprises the state of the dynamical system. We further note that the dynamical system includes a nonlinear component, which arises from the (sub)gradient evaluations ∇f and ∂g , and the system is subjected to external disturbances e_k that arise from the event-based communication. The detailed derivation and the corresponding matrices for the dynamics in Fig. 2 are included in App. C. Our convergence analysis will build on the dynamical systems model of Alg. 2. While our analysis is inspired by earlier works, such as (Nishihara et al., 2015) and (Lessard et al., 2016), the Lyapunov function that is used to prove convergence rates are different due to the external disturbances caused by the event-based communication.

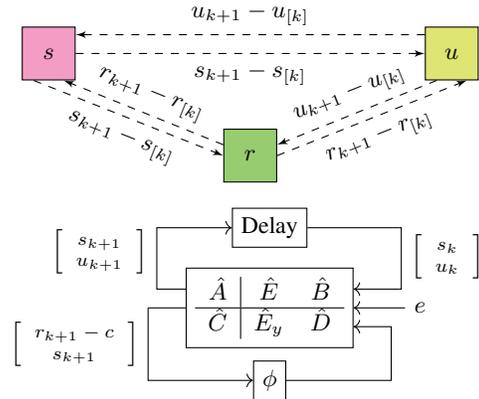


Figure 2: The figure visualizes the event-based communication structure of Alg. 2 at the top and a discrete-time dynamical system which represents the sequence generated by the event-based ADMM algorithm on the bottom. The function ϕ is nonlinear and represents the evaluation of (sub)gradients.

4. Convergence Analysis

This section provides convergence guarantees for the event-based learning algorithm (Alg. 2). The detailed proof for Thm. 4.1 is provided in App. D.

Theorem 4.1. *Let Assumption 3.1 and 3.2 be satisfied and let the step-size for Alg. 2 be $\rho = \kappa^\epsilon \sqrt{mL}/(\underline{\sigma}(A)\bar{\sigma}(A))$, for some $\epsilon \geq 0$ and $\alpha \in (0.675, 1 + \sqrt{1-1/\sqrt{\kappa}})$, $\kappa =$*

$L\bar{\sigma}^2(A)/(m\sigma^2(A))$, where σ and $\bar{\sigma}$ denote the minimum and maximum singular value of a matrix, respectively. Then, for large enough κ , the following bound holds:

$$|\xi_k - \xi_*|^2 \leq \kappa_P |\xi_0 - \xi_*|^2 \left(1 - \frac{\alpha}{4\kappa^{\epsilon + \frac{1}{2}}}\right)^{2k} + \frac{60\kappa^{2+2\epsilon}}{\alpha(1-|\alpha-1|)} \Delta^2,$$

with $\xi_k = (s_k, u_k)$, and where $s_k = Bz_k$, u_k is the dual variable, and ξ_* the optimizer corresponding to (3). Furthermore, $\Delta = \Delta^r + \Delta^s + \Delta^u + T(\bar{\chi}^r + \bar{\chi}^s + \bar{\chi}^u)$ represents the error arising from the event-based communication and $\kappa_P = (2\sqrt{\kappa} - 1 + \sqrt{4\kappa(\alpha-1)^2 + 1}) / (2\sqrt{\kappa} - 1 - \sqrt{4\kappa(\alpha-1)^2 + 1})$.

We conclude the section by highlighting a few important points.

(i) For $\epsilon = 0$, $\alpha = 1$, the bound is considerably simplified to

$$|\xi_k - \xi_*|^2 \leq 2|\xi_0 - \xi_*|^2 \left(1 - \frac{1}{4\sqrt{\kappa}}\right)^{2k} + 60\kappa^2 \Delta^2,$$

which shows that the convergence rate scales with $1/\sqrt{\kappa}$ and is therefore accelerated. This also highlights that the same convergence rate (up to constants) can be achieved with the event-based learning algorithm stated in Alg. 1 compared to a standard ADMM algorithm. As we will show in the numerical experiments, our event-based algorithm reduces communication without any significant reduction in accuracy.

(ii) The bound from Thm. 4.1 also highlights how the communication thresholds Δ affect the solution accuracy. In the simplified scenario with $\epsilon = 0$, $\alpha = 1$ (the more general scenario follows the same rationale), the solution accuracy is bounded by $|\xi_k - \xi_*| \leq 8\kappa\Delta$, for large enough k . This means that the solution accuracy of Alg. 2 is proportional to the condition number κ and Δ .

(iii) We can therefore easily ensure convergence, by choosing a time-varying $\Delta = \Delta_k$ such that $\Delta_k \rightarrow 0$. The formal statement is included and derived in App. F. We also obtain precise nonasymptotic bounds. For example, if $\Delta_k = \Delta_0/(k+1)^t$ for any $t > 0$, we conclude that the error converges with $\mathcal{O}(1/k^t)$ (see again App. F).

(iv) If f fails to be strongly convex, we can include a small regularizer, for example of the type $m|x|^2/2$. Choosing a diminishing regularizer with $m = \mathcal{O}(1/k^2)$ and a diminishing threshold $\Delta_k = \mathcal{O}(1/k^4)$ can be shown to result in an accelerated convergence rate of $\mathcal{O}(1/k^2)$.

(v) The topology of the communication network, represented by the matrix A , directly influences the convergence rate, through the condition number $\kappa = L\bar{\sigma}^2(A)/(m\sigma^2(A))$. This formulation allows us to generalize our convergence results beyond simple client-server architectures. See App. A.2 for a detailed discussion on how agent network topology is encoded in the matrix A .

5. Numerical Experiments

This section discusses the performance of Alg. 1 in numerical experiments, highlighting that Alg. 1 achieves fast convergence while reducing communication. Numerical experiments with the more general version (Alg. 2) are included in App. G, where distributed training over a network of agents is explored. We also investigate the trade-off between communication load and solution accuracy achieved by selecting different communication thresholds. The communication load is calculated by counting the number of triggered communications for T_{\max} number of steps and normalizing according to the full communication case of one data package per round.

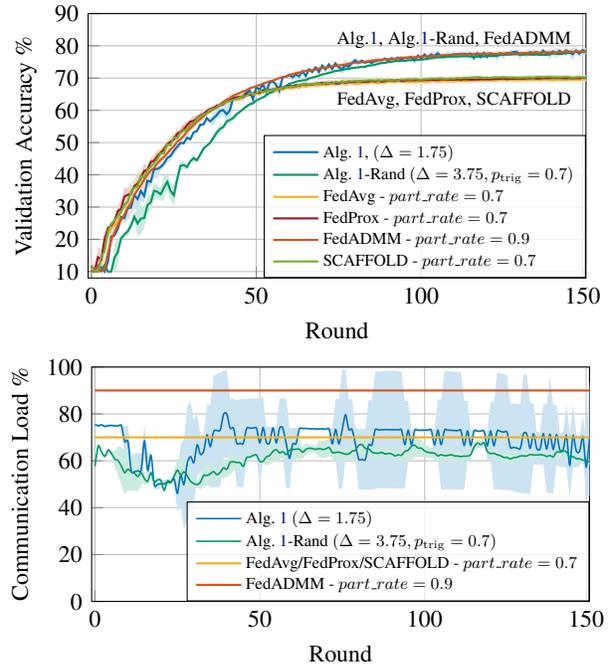


Figure 3: Validation accuracy (top) and communication load percentage (bottom) over 150 communication rounds for training a CIFAR-10 classifier. The results indicate that Alg. 1 achieves top accuracy at a lower communication rate. The plots compare the performance of various algorithms (Vanilla and randomized), FedAvg, FedProx, FedADMM, and SCAFFOLD. Notably, ADMM-based methods (Alg. 1, Alg. 1-Rand and FedADMM) demonstrate better convergence by reaching up to 78% test accuracy, compared to other algorithms FedAvg, FedProx and SCAFFOLD, which reach only 70% accuracy. Among ADMM-based methods, Alg. 1 and Alg. 1-Rand achieve the same accuracy with over 20% less communication load. Communication load curves are smoothed using a window length of three for visualization purposes.

Algorithm	MNIST Target Accuracy			CIFAR-10 Target Accuracy			
	80%	85%	90%	70%	75%	77%	78%
Alg. 1 - Randomized	629	693	<u>1723</u>	12531	13422	<u>15008</u>	18376
Alg. 1 - Vanilla	816	1285	1710	12214	<u>14780</u>	14780	<u>20690</u>
FedADMM (Zhou & Li, 2023)	<u>800</u>	<u>1200</u>	>2000	12000	15000	21000	27000
FedAvg (McMahan et al., 2017)	<u>800</u>	2000	N/A	3000	N/A	N/A	N/A
FedProx (Li et al., 2020a)	1000	2000	N/A	<u>6000</u>	N/A	N/A	N/A
SCAFFOLD (Karimireddy et al., 2020)	1600	2000	3200	12000	N/A	N/A	N/A

Table 1: The total number of communication events required by each algorithm to achieve the target accuracies for the MNIST and CIFAR-10 classifiers within 100 and 150 rounds, respectively. “N/A” indicates cases where the target accuracy was not reached within the specified rounds. Parameter choices for the algorithms are detailed in Appendix G. Reported values are averages over multiple experiments with different random seeds, with standard deviations below 2%, making them negligible. The corresponding communication load and accuracy trends for the rightmost column of CIFAR-10 are shown in Fig. 3. The results emphasize the trade-off between achieving higher validation accuracy and maintaining communication efficiency across various algorithm configurations.

Due to space limitations, we present two examples in this section. Further experiments (including linear regression, LASSO, and distributed training over a network of agents) are presented in App. G. App. G also includes hyperparameters for the experiments, model details and discusses the effect of communication failures.

We start by evaluating the performance of Alg. 1 on MNIST (Deng, 2012) and CIFAR-10 (Krizhevsky, 2009). Tab. 1 reports the total number of communication events required by each algorithm to achieve the target accuracies for the MNIST and CIFAR-10 classifiers. Our event-based algorithm consistently requires fewer communication events to achieve high accuracies compared to baseline methods. This reduction is attributed to the selective triggering mechanism, which prevents unnecessary communication while ensuring convergence. For instance, on the CIFAR-10 dataset, our approach achieved 78% accuracy with a cost of 18,376 communication events, compared to 27,000 for FedADMM.

The comparison with other federated learning methods emphasizes the challenges associated with non-i.i.d. data distribution and communication overhead. FedAvg, as highlighted in (Li et al., 2020c; Glasgow et al., 2022), experiences slowdowns in the presence of non-i.i.d. data, and increasing participation does not necessarily alleviate this issue. FedProx has the same issue and is unable to converge to a classifier that generalizes across all digits. FedADMM and SCAFFOLD can indeed cope with non-i.i.d. data, in general, both achieving high classification accuracies. However, FedADMM has disadvantages arising from the random sampling mechanism and SCAFFOLD suffers from an additional communication load to communicate two variables (client drift and local model). Notably, all baselines employ a random selection of agents, which, in non-i.i.d. scenarios, misses crucial changes and results in a waste of communication resources. Our method addresses these challenges by

adopting an event-based agent selection approach and outperforms all baselines by yielding uniformly better trade-off curves.

6. Conclusion

We introduce an event-based distributed learning approach that effectively reduces communication overhead by triggering events only when local models undergo significant changes. The method, based on over-relaxed ADMM, exhibits accelerated convergence rates in convex settings, demonstrates robustness to communication failures, and outperforms common baselines such as FedAvg, FedProx, SCAFFOLD and FedADMM in our experiments, which include an MNIST and CIFAR-10 learning task. The experiments highlight that savings of more than 35% are possible without significantly degrading the solution accuracy (less than 1%). Our method allows for explicit trade-offs between communication load and solution accuracy, making it promising for large-scale learning systems with heterogeneous data and communication constraints.

Limitations: While our approach offers significant improvements in communication efficiency, it has not yet accounted for adversarial attacks, such as gradient poisoning, or unreliable nodes in the network. These factors could potentially degrade the robustness of the method. Nevertheless, our event-based methodology can be integrated with robust aggregation or anomaly detection methods (Pillutla et al., 2022; Yin et al., 2018) to improve security without compromising communication efficiency. Additionally, this method has not been specifically analyzed with respect to differential privacy and does not address privacy concerns in the current formulation.

Discussion: In this article, we focused on communication efficiency and convergence relationship under some con-

straints. However, practical implementations often face additional challenges, such as limited bandwidth, high latency, and network failures. Although we have not explicitly named these factors, they can be effectively modeled through our packet drop framework, which provides a foundation for handling such constraints.

Furthermore, while we initially framed our algorithm as a synchronous method, event-based communication can also be adapted to function in asynchronous systems. This flexibility allows the method to accommodate real-world scenarios with varying degrees of network reliability and synchronization.

Finally, our method is compatible with compression and quantization techniques, which can further reduce the size of the models exchanged during communication events and improve communication efficiency. This can also help minimize the amount of data stored in the agents, contributing to more efficient memory usage and reducing the overall communication load.

Impact Statement

This article advances machine learning by introducing a communication-efficient, event-based distributed learning method, enabling scalability for resource-constrained systems. While the work does not explicitly address adversarial robustness or differential privacy, future extensions could explore these aspects to further improve security in applications. We see no immediate negative societal impacts.

References

- Acar, D. A. E., Zhao, Y., Navarro, R. M., Mattina, M., Whatmough, P. N., and Saligrama, V. Federated learning based on dynamic regularization. *Proceedings of the International Conference on Learning Representations*, pp. 1–36, 2021.
- Asad, M., Shaukat, S., Hu, D., Wang, Z., Javanmardi, E., Nakazato, J., and Tsukada, M. Limitations and future aspects of communication costs in federated learning: A survey. *Sensors*, 23(17):7358, 2023.
- Bastianello, N., Carli, R., Schenato, L., and Todescato, M. Asynchronous distributed optimization over lossy networks via relaxed ADMM: Stability and linear convergence. *IEEE Transactions on Automatic Control*, 66(6): 2620–2635, 2021.
- Bertsekas, D. P. and Tsitsiklis, J. N. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, 1989.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. Alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2010.
- Brunzema, P., von Rohr, A., Solowjow, F., and Trimpe, S. Event-triggered time-varying Bayesian optimization. *Transactions on Machine Learning Research*, 2025.
- Candès, E. J., Li, X., Ma, Y., and Wright, J. Robust principal component analysis. *Journal of the ACM*, 58(3):1–37, 2011.
- Cao, X., Başar, T., Diggavi, S., Eldar, Y. C., Letaief, K. B., Poor, H. V., and Zhang, J. Communication-efficient distributed learning: An overview. *Journal on Selected Areas in Communications*, 41(4):851–873, 2023.
- Chen, T., Giannakis, G. B., Sun, T., and Yin, W. LAG: Lazily aggregated gradient for communication-efficient distributed learning. *Proceedings of the International Conference on Advances in Neural Information Processing Systems*, pp. 5055–5065, 2018.
- Deng, L. The MNIST database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- Elgabli, A., Park, J., Bedi, A. S., Bennis, M., and Aggarwal, V. GADMM: Fast and communication efficient framework for distributed machine learning. *Journal of Machine Learning Research*, 21:1–39, 2020.
- Gao, L., Fu, H., Li, L., Chen, Y., Xu, M., and Xu, C.-Z. FedDC: Federated learning with non-iid data via local drift decoupling and correction. *Proceedings of the Conference on Computer Vision and Pattern Recognition*, 2022.
- Garofalakis, M. and Samoladas, V. Distributed query monitoring through convex analysis: Towards composable safe zones. *Proceedings of the International Conference on Database Theory*, 68(14):1–18, 2017.
- Ghadikolaei, H. S., Stich, S. U., and Jaggi, M. LENA: Communication-efficient distributed learning with self-triggered gradient uploads. *Proceedings of the International Conference on Artificial Intelligence and Statistics*, pp. 3943–3951, 2021.
- Glasgow, M., Yuan, H., and Ma, T. Sharp bounds for federated averaging (local sgd) and continuous perspective. *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 151, 2022.
- Gong, Y., Li, Y., and Freris, N. M. FedADMM: A robust federated deep learning framework with adaptivity to system heterogeneity. *Proceedings of the IEEE International Conference on Data Engineering*, pp. 2575–2587, 2022.

- He, S., Zheng, J., Feng, M., and Chen, Y. Communication-efficient federated learning with adaptive consensus ADMM. *Applied Sciences*, 13(9):5270, 2023.
- Hegazy, M., Leluc, R., Ting Li, C., and Dieuleveut, A. Compression with exact error distribution for federated learning. *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 238:613–621, 2024.
- Hendrikx, H., Xiao, L., Bubeck, S., Bach, F., and Mas-soulié, L. Statistically preconditioned accelerated gradient method for distributed optimization. *Proceedings of the International Conference on Machine Learning*, 119: 4203–4227, 2020.
- Kairouz et al., P. Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1-2):1–210, 2021.
- Karimireddy, S. P., Kale, S., Mohri, M., Reddi, S. J., Stich, S. U., and Suresh, A. T. SCAFFOLD: Stochastic controlled averaging for federated learning. *Proceedings of International Conference on Machine Learning*, 119: 5132–5143, 2020.
- Kovalev, D., Salim, A., and Richtárik, P. Optimal and practical algorithms for smooth and strongly convex decentralized optimization. *Proceedings of the Conference on Advances in Neural Information Processing Systems*, pp. 18342–18352, 2020.
- Krizhevsky, A. Learning multiple layers of features from tiny images. *Technical Report, University of Toronto*, 2009.
- Lessard, L., Recht, B., and Packard, A. Analysis and design of optimization algorithms via integral quadratic constraints. *SIAM Journal on Optimization*, 26(1):57–95, 2016.
- Li, M., Sanjabi, M., and Jaggi, M. Federated optimization in heterogeneous networks. *Proceedings of the Conference on Machine Learning and Systems*, 2:429–450, 2020a.
- Li, T., Sahu, A. K., Talwalkar, A., and Smith, V. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020b.
- Li, X., Yang, W., Huang, K., Wang, S., and Zhang, Z. On the convergence of FedAvg on non-i.i.d. *Proceedings of the International Conference on Learning Representations*, pp. 1–26, 2020c.
- Liu, Y., Xu, W., Wu, G., Tian, Z., and Ling, Q. Communication-censored ADMM for decentralized consensus optimization. *IEEE Transactions on Signal Processing*, 67(10):2565–2579, 2019.
- Liu, Y., Sun, Y., and Yin, W. Decentralized learning with lazy and approximate dual gradients. *IEEE Transactions on Signal Processing*, 69:1362–1377, 2021.
- Mao, Y., Zhao, Z., Yan, G., Liu, Y., Lan, T., Song, L., and Ding, W. Communication-efficient federated learning with adaptive quantization. *ACM Transactions on Intelligent Systems and Technology*, 13(4), 2022.
- McMahan, H. B., Moore, E., Ramage, D., and Hampson, S. Communication-efficient learning of deep networks from decentralized data. *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 54: 1273–1282, 2017.
- Miskowicz, M. Send-On-Delta Concept: An event-based data reporting strategy. *Sensors*, 6(1):49–63, 2006.
- Muehlebach, M. and Jordan, M. I. A dynamical systems perspective on Nesterov acceleration. *Proceedings of the International Conference on Machine Learning*, 97: 4656–4662, 2019.
- Muehlebach, M. and Jordan, M. I. Continuous-time lower bounds for gradient-based algorithms. *Proceedings of the International Conference on Machine Learning*, 119: 7088–7096, 2020.
- Nabli, A. and Oyallon, E. DADAO: Decoupled accelerated decentralized asynchronous optimization. *Proceedings of the International Conference on Machine Learning*, 202: 25604–25626, 2023.
- Nedic, A., Olshevsky, A., and Rabbat, M. G. Network topology and communication-computation tradeoffs in decentralized optimization. *Proceedings of the IEEE*, 106 (5):953–976, 2018.
- Nishihara, R., Lessard, L., Recht, B., Packard, A., and Jordan, M. I. A general analysis of the convergence of ADMM. *Proceedings of the International Conference on Machine Learning*, 37:343–352, 2015.
- Pillutla, K., Kakade, S. M., and Harchaoui, Z. Robust aggregation for federated learning. *IEEE Transactions on Signal Processing*, 70:1142–1154, 2022.
- Polyak, B. T. *Introduction to optimization*. Optimization Software Inc., Publications Division, New York, 1987.
- Qiu, Y., Lei, Y., and Wang, G. PSRA-HGADMM: A communication efficient distributed ADMM algorithm. *Proceedings of the 52nd International Conference on Parallel Processing*, pp. 82–91, 2023.
- Reisizadeh, A., Jadbabaie, A., Mokhtari, A., Hassani, H., and Pedarsani, R. FedPAQ: A communication-efficient federated learning method with periodic averaging and

- quantization. *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 108:2021–2031, 2020.
- Shamir, O., Srebro, N., and Zhang, T. Communication-efficient distributed optimization using an approximate Newton-type method. *Proceedings of the International Conference on Machine Learning*, 32(2):1000–1008, 2014.
- Shi, Y., Zhang, Y., Zhang, P., Xiao, Y., and Niu, L. Federated learning with ℓ_1 regularization. *Pattern Recognition Letters*, 172:15–21, 2023.
- Shokri, R. and Shmatikov, V. Privacy-preserving deep learning. *Proceedings of the SIGSAC Conference on Computer and Communications Security*, pp. 1310–1321, 2015.
- Singh, N., Data, D., George, J., and Diggavi, S. SPARQ-SGD: Event-triggered and compressed communication in decentralized optimization. *IEEE Transactions on Automatic Control*, 68(2):721–736, 2023.
- Solowjow, F. and Trimpe, S. Event-triggered learning. *Automatica*, 117:109009, 2020.
- Song, Y., Wang, Z., and Zuazua, E. FedADMM-InSa: An inexact and self-adaptive admm for federated learning. *Neural Networks*, 181:106772, 2025.
- Su, W., Boyd, S., and Candes, E. J. A differential equation for modeling Nesterov’s accelerated gradient method: Theory and insights. *Journal of Machine Learning Research*, 17(153):1–43, 2016.
- Tong, G. and Muehlebach, M. A dynamical systems perspective on discrete optimization. *Proceedings of Machine Learning Research*, 211:1–14, 2023.
- Tran Dinh, Q., Pham, N. H., Phan, D., and Nguyen, L. FedDR—randomized Douglas-Rachford splitting algorithms for nonconvex federated composite optimization. *Proceedings of the International Conference on Advances in Neural Information Processing Systems*, 34:30326–30338, 2021.
- Umlauft, J. and Hirche, S. Feedback linearization based on gaussian processes with event-triggered online learning. *IEEE Transactions on Automatic Control*, 65(10):4154–4169, 2019.
- Wang, G., Wang, D., Li, C., and Lei, Y. The fast inertial ADMM optimization framework for distributed machine learning. *Future Generation Computer Systems*, 164:107575, 2025.
- Wang, H., Sievert, S., Liu, S., Charles, Z., Papailiopoulos, D., and Wright, S. ATOMO: Communication-efficient learning via atomic sparsification. *Proceedings of the International Conference on Advances in Neural Information Processing Systems*, pp. 9850–9861, 2018.
- Wang, H., Marella, S., and Anderson, J. FedADMM: A federated primal-dual algorithm allowing partial participation. *Proceedings of the IEEE Conference on Decision and Control*, pp. 287–294, 2022.
- Wei, E. and Ozdaglar, A. Distributed alternating direction method of multipliers. *Proceedings of the IEEE Conference on Decision and Control*, pp. 5445–5450, 2012.
- Wei Liu, Li Chen, and Wenyi Zhang. Decentralized federated learning: Balancing communication and computing costs. *IEEE Transactions on Signal and Information Processing over Networks*, 8:131–143, 2021.
- Wibisono, A., Wilson, A. C., and Jordan, M. I. A variational perspective on accelerated methods in optimization. *Proceedings of the National Academy of Sciences*, 113(47):E7351–E7358, 2016.
- Yin, D., Chen, Y., Kannan, R., and Bartlett, P. Byzantine-robust distributed learning: Towards optimal statistical rates. *Proceedings of the International Conference on Machine Learning*, 80:5650–5659, 2018.
- Yu, Z. and Freris, N. M. Communication-efficient distributed optimization with adaptability to system heterogeneity. *Proceedings of the IEEE Conference on Decision and Control*, pp. 3321–3326, 2023.
- Zhang, X., Hong, M., Dhople, S., Yin, W., and Liu, Y. FedPD: A federated learning framework with adaptivity to non-iid data. *IEEE Transactions on Signal Processing*, 69:6055–6070, 2021.
- Zhang, Z., Yang, S., and Xu, W. Decentralized ADMM with compressed and event-triggered communication. *Neural Networks*, 165:472–482, 2023.
- Zhang, Z., Yang, S., Xu, W., and Di, K. Privacy-preserving distributed ADMM with event-triggered communication. *IEEE Transactions on Neural Networks and Learning Systems*, 35(2):2835–2847, 2024.
- Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., and Chandra, V. Federated learning with non-iid data. *arxiv:1806.00582*, 2018.
- Zheng, S., Ye, T., Li, X., and Gao, M. Federated learning via consensus mechanism on heterogeneous data: A new perspective on convergence. *IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 7595–7599, 2024.

Zhou, S. and Li, G. Y. Federated learning via inexact ADMM. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8):9699–9708, 2023.

Zhu, H. and Ling, Q. Broadcast: Reducing both stochastic and compression noise to robustify communication-efficient federated learning. *arXiv:2104.06685*, 2021.

A. Communication Structure

This section discusses the sharing problem and consensus reaching over graphs as two special cases of the more general constrained minimization problem

$$\min_{x \in \mathbb{R}^p, z \in \mathbb{R}^q} f(x) + g(z), \quad \text{subject to } Ax + Bz = c, \quad (4)$$

with variables $x \in \mathbb{R}^p$ and $z \in \mathbb{R}^q$ and constant matrices $A \in \mathbb{R}^{r \times p}$, $B \in \mathbb{R}^{r \times q}$, and $c \in \mathbb{R}^r$. The objective function is decomposed into a smooth part $f: \mathbb{R}^p \rightarrow \mathbb{R}$ and nonsmooth part $g: \mathbb{R}^q \rightarrow \mathbb{R}$. The communication structure of the problem formulation (4) is shown in Fig. 4, where primal, dual and auxiliary variables are treated as different communication nodes.

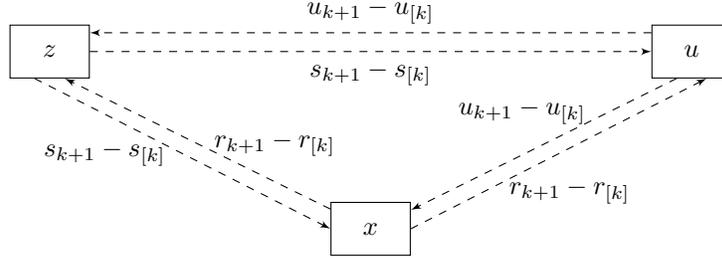


Figure 4: The communication structure that arises from Alg. 2, where $s := Bz$, $r := Ax$, and u denotes the dual variable.

A.1. Sharing Problem

We will show that the event-based communication structure introduced in Fig. 4 simplifies considerably for the sharing problem. The sharing problem takes the following form,

$$\min_{x^1, \dots, x^N \in \mathbb{R}^p} \sum_{i=1}^N f^i(x^i) + g\left(\sum_{i=1}^N x^i\right),$$

and arises as a special case from (4) when choosing $f(x) = \sum_{i=1}^N f^i(x^i)$, $x = (x^1, x^2, \dots, x^N) \in \mathbb{R}^{Np}$, $A = I_{Np}$, $B = -(I_p, I_p, \dots, I_p)$, $c = 0$. The problem can be solved via the following updates, by agents $i = 1, \dots, N$:

$$x_{k+1}^i = \operatorname{argmin}_{x^i \in \mathbb{R}^p} f^i(x^i) + \frac{\rho}{2} \left| x^i - x_k^i + \hat{h}_k \right|^2, \quad (5)$$

and by agent $N + 1$:

$$\begin{aligned} \bar{x}_{k+1} &= \frac{1}{N} \sum_{i=1}^N \hat{x}_{k+1}^i \\ z^{k+1} &= \operatorname{argmin}_{z \in \mathbb{R}^p} g(Nz) + \frac{N\rho}{2} \left| z - \bar{x}_{k+1} - \frac{1}{\rho} u^k \right|^2 \\ u_{k+1} &= u_k + \rho (\bar{x}_{k+1} - z_{k+1}) \\ h_{k+1} &= \bar{x}_{k+1} - z_{k+1} + \frac{1}{\rho} u_{k+1}. \end{aligned} \quad (6)$$

For the sharing problem, the general communication scheme in Fig. 4 reduces to the diagram in Fig. 5, where each node communicates their local variable in an event-based manner.

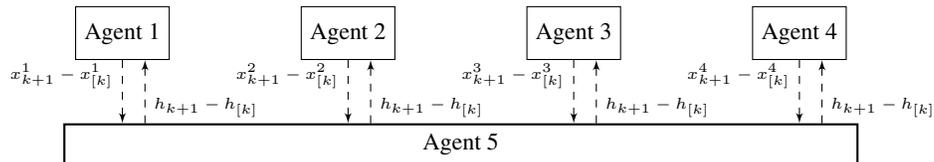


Figure 5: The diagram visualizes the communication structure for the sharing problem for $N = 4$ agents.

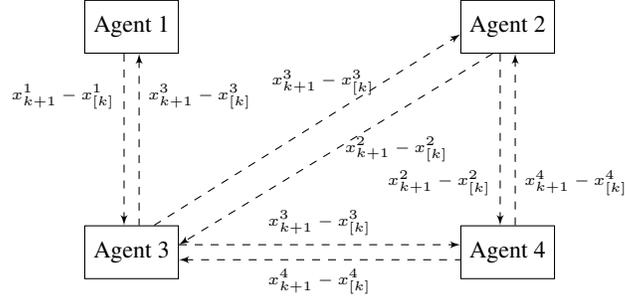


Figure 6: The diagram visualizes the communication structure for a distributed learning problem over a graph that connects four agents with four edges.

A.2. Consensus over a Graph

As another example, we will show that (4) also generalizes to distributed learning scenarios over graphs. We consider a network topology, captured by an undirected connected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{1, \dots, N\}$ is the set of vertices and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges. Each agent (vertex) has a local data distribution, and the aim is to train a model without a central server to aggregate the collected information. The problem can be formulated as follows:

$$\min_{x^i \in \mathbb{R}^p, z^{ij} \in \mathbb{R}^p} \sum_{i=1}^N f^i(x^i), \quad \text{subject to } x^i = z^{ij}, x^j = z^{ij}, \quad \forall (i, j) \in \mathcal{E}.$$

Similar to the formulation in (Yu & Freris, 2023), we define transmitter and receiver matrices $\hat{A}_t, \hat{A}_r \in \mathbb{R}^{|\mathcal{E}| \times N}$ for all edges, i.e.,

$$\left[\hat{A}_t \right]_{ei} = \left[\hat{A}_r \right]_{ej} = \begin{cases} 1 & (i, j) \in \mathcal{E} \\ 0 & \text{otherwise} \end{cases}, \quad \forall e \in \mathcal{E}.$$

By stacking $x^i, z^{ij} \in \mathbb{R}^p$ into column vectors $x \in \mathbb{R}^{Np}, z \in \mathbb{R}^{|\mathcal{E}|p}$, respectively, we conclude that distributed learning over graphs is indeed a special case of (4),

$$\min_{x \in \mathbb{R}^{Np}, z \in \mathbb{R}^{|\mathcal{E}|p}} f(x), \quad \text{subject to } \begin{bmatrix} \hat{A}_t \otimes I_p \\ \hat{A}_r \otimes I_p \end{bmatrix} x = \begin{bmatrix} I_{|\mathcal{E}|p} \\ I_{|\mathcal{E}|p} \end{bmatrix} z,$$

where \otimes denotes the Kronecker product and I_p the identity matrix. Thus, the matrices A and B encode the topology of the communication graph, which will affect the convergence rates as highlighted with our main result Thm. 4.1 where the convergence rate is dictated by the value $\kappa = \bar{\sigma}(A)L/(\underline{\sigma}(A)m)$.

The resulting instance of Alg. 1 takes the following form:

$$\begin{aligned} x_{k+1}^i &= \operatorname{argmin}_{x^i \in \mathbb{R}^p} f_i(x_i) + \frac{|\mathcal{N}_i| \rho}{2} \left| x^i - \frac{1}{2} (x_i^k - \bar{x}_k^i) + \frac{1}{\rho} p_k^i \right|^2 \\ \bar{x}_{k+1}^i &= \frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} \hat{x}_{k+1}^j \\ p_{k+1}^i &= p_k^i + \frac{\rho}{2} (x_{k+1}^i - \bar{x}_{k+1}^i), \end{aligned} \tag{7}$$

where \mathcal{N}_i represents the set containing the neighbors of the agent i and $|\mathcal{N}_i|$ is the number of vertices. In the event based-communication setting, an agent transmits its local model (x_{k+1}^i) to the neighbors only if there has been a significant change in the local model. Fig. 6 shows an example with four agents, each communicating local variables only.

B. Proof of Thm. 2.3

Proof. We will establish the convergence rate by analyzing the behavior of a carefully chosen Lyapunov function. Let us begin by formulating the augmented Lagrangian for (1):

$$\mathcal{L}_\rho(x, z, y) = \sum_{i=1}^N f^i(x^i) + g(z) + \sum_{i=1}^N (y^i)^\top (x^i - z) + \frac{\rho}{2} \sum_{i=1}^N |x^i - z|_2^2, \quad (8)$$

where $x = (x^1, \dots, x^N)$ and $y = (y^1, \dots, y^N)$ represent the primal and dual variables respectively, and $\rho > 0$ is the penalty parameter.

We express the ADMM updates in Alg. 1 for $\alpha = 1$ as follows by introducing the scaled dual variable $u^i = \frac{1}{\rho} y^i$:

$$x_{k+1}^i = \arg \min_{x^i} \left(f^i(x^i) + \frac{\rho}{2} |x^i - \hat{z}_k + u_k^i|_2^2 \right), \quad \forall i \in \{1, \dots, N\}, \quad (9)$$

$$z_{k+1} = \arg \min_z \left(g(z) + \frac{N\rho}{2} \sum_{i=1}^N |\hat{x}_{k+1}^i - z + \hat{u}_k^i|_2^2 \right), \quad (10)$$

$$u_{k+1}^i = u_k^i + x_{k+1}^i - \hat{z}_{k+1}, \quad \forall i \in \{1, \dots, N\}. \quad (11)$$

Here, we use the notation $\hat{z}_{k+1} = z_{k+1} + \varepsilon_{k+1}^z$, $\hat{x}_{k+1}^i = x_{k+1}^i + \varepsilon_{k+1}^{x,i}$, and $\hat{u}_{k+1}^i = u_{k+1}^i + \varepsilon_{k+1}^{u,i}$ to account for errors emerging from event-based communication.

From (9) and (10), we derive the following first-order optimality conditions for x_{k+1}^i and z_{k+1}

$$0 = \nabla f^i(x_{k+1}^i) + \rho(x_{k+1}^i - \hat{z}_k + u_k^i) \quad (12)$$

$$0 \in \partial g(z_{k+1}) + \rho \sum_{i=1}^N (z_{k+1} - \hat{x}_{k+1}^i - \hat{u}_k^i). \quad (13)$$

We then define the Lyapunov function:

$$V_k = |z_k - z_*|_2^2 + \frac{1}{N} \sum_{i=1}^N |u_k^i - u_*^i|_2^2, \quad (14)$$

where (u_*^i, z_*) denotes the optimal dual and consensus variables. Our goal is to demonstrate that this Lyapunov function is monotonically decreasing.

The optimality condition in (12) implies that x_{k+1}^i minimizes,

$$f^i(x) + \rho (u_{k+1}^i + z_{k+1} - z_k)^\top x + \rho (\varepsilon_{k+1}^z - \varepsilon_k^z)^\top x.$$

From this minimization, we can derive the following inequality,

$$f^i(x_{k+1}^i) - f^i(x_*^i) \leq \rho (u_{k+1}^i + z_{k+1} - z_k)^\top (x_*^i - x_{k+1}^i) + \rho (\varepsilon_{k+1}^z - \varepsilon_k^z)^\top (x_*^i - x_{k+1}^i). \quad (15)$$

Similarly, the optimality condition (13) indicates that z_{k+1} minimizes,

$$g(z) - \rho \sum_{i=1}^N (u_{k+1}^i + \varepsilon_{k+1}^{d,i} + \varepsilon_{k+1}^z)^\top z, \quad (16)$$

where $\varepsilon_{k+1}^{d,i} = \varepsilon_{k+1}^{x,i} + \varepsilon_k^{u,i}$. This minimization leads to,

$$g(z_{k+1}) - g(z_*) \leq \rho \sum_{i=1}^N (u_{k+1}^i)^\top (z_{k+1} - z_*) - \rho \sum_{i=1}^N (\varepsilon_{k+1}^{d,i} + \varepsilon_{k+1}^z)^\top (z_{k+1} - z_*). \quad (17)$$

By adding (17) and the sum over i of (15), and applying the conditions $x_*^i - z_* = 0$ along with the relation $x_*^i - x_{k+1}^i = -r_{k+1}^i - (z_{k+1} - z_*)$, we obtain,

$$g(z_{k+1}) - g(z_*) + \sum_{i=1}^N (f^i(x_{k+1}^i) - f^i(x_*^i)) \leq -\rho \sum_{i=1}^N (u_{k+1}^i)^\top r_{k+1}^i - \rho \sum_{i=1}^N (z_{k+1} - z_k)^\top (r_{k+1}^i + (z_{k+1} - z_*)) + \rho \sum_{i=1}^N (\varepsilon_{k+1}^z - \varepsilon_k^z)^\top (x_*^i - x_{k+1}^i) + \rho \sum_{i=1}^N (\varepsilon_{k+1}^{d,i} + \varepsilon_{k+1}^z)^\top (z_{k+1} - z_*), \quad (18)$$

where the residual is defined as $r_k^i := x_k^i - z_k$.

Since (x_*^i, z_*, u_*^i) is a saddle point of \mathcal{L}_0 in (8), i.e., $\mathcal{L}_0(x_*, z_*, u_*) \leq \mathcal{L}_0(x_{k+1}, z_{k+1}, u_*)$, we have,

$$-g(z_{k+1}) + g(z_*) - \sum_{i=1}^N (f^i(x_{k+1}^i) - f^i(x_*^i)) \leq \sum_{i=1}^N u_*^i \top (x_{k+1}^i - z_{k+1}) = \sum_{i=1}^N u_*^i \top r_{k+1}^i. \quad (19)$$

By adding (18) and (19) and multiplying by $\frac{2}{\rho}$, we arrive at

$$0 \geq 2 \underbrace{\sum_{i=1}^N (u_{k+1}^i - u_*^i)^\top r_{k+1}^i}_{(I)} + 2 \underbrace{\sum_{i=1}^N (z_{k+1} - z_k)^\top (r_{k+1}^i + (z_{k+1} - z_*))}_{(II)} - 2 \sum_{i=1}^N (\varepsilon_{k+1}^z - \varepsilon_k^z)^\top (x_*^i - x_{k+1}^i) - 2 \sum_{i=1}^N (\varepsilon_{k+1}^{d,i} + \varepsilon_{k+1}^z)^\top (z_{k+1} - z_*). \quad (20)$$

Here $(u_{k+1}^i - u_*^i)^\top r_{k+1}^i$ can be written as $(u_{k+1}^i - u_*^i)^\top (u_{k+1}^i - u_k^i + \varepsilon_{k+1}^z)$ which splits into

$$(u_{k+1}^i - u_*^i)^\top r_{k+1}^i = \frac{1}{2} (u_{k+1}^i - u_k^i)^\top \varepsilon_{k+1}^z + \frac{1}{2} |u_{k+1}^i - u_k^i|^2 + \frac{1}{2} (u_{k+1}^i - u_k^i)^\top (u_{k+1}^i - u_k^i + \varepsilon_{k+1}^z) + (u_k^i - u_*^i)^\top (u_{k+1}^i - u_*^i) + (u_k^i - u_*^i)^\top \varepsilon_{k+1}^z - |u_k^i - u_*^i|^2.$$

Next, we substitute $u_{k+1}^i - u_k^i = r_{k+1}^i - \varepsilon_{k+1}^z$ and use the following squared norm identity,

$$\frac{1}{2} |u_{k+1}^i - u_k^i|^2 + (u_k^i - u_*^i)^\top (u_{k+1}^i - u_*^i) = \frac{1}{2} |u_{k+1}^i - u_*^i|^2 + \frac{1}{2} |u_k^i - u_*^i|^2.$$

Consequently, we can expand terms (I) and (II) as follows:

$$(I) = \sum_{i=1}^N (|u_{k+1}^i - u_*^i|^2 - |u_k^i - u_*^i|^2 + 2(u_{k+1}^i - u_*^i)^\top \varepsilon_{k+1}^z + |\varepsilon_{k+1}^z|^2 + |r_{k+1}^i|^2 - 2(\varepsilon_{k+1}^z)^\top r_{k+1}^i) \\ (II) = \sum_{i=1}^N (|r_{k+1}^i + (z_{k+1} - z_k)|^2 + |z_{k+1} - z_*|^2 - |z_k - z_*|^2 - |r_{k+1}^i|^2).$$

Substituting these expansions back into (20) and using the definition of our Lyapunov function from (14), we can express the decrease in the Lyapunov function as,

$$0 \geq N(V_{k+1} - V_k) + \sum_{i=1}^N (2(u_{k+1}^i - u_*^i)^\top \varepsilon_{k+1}^z + |\varepsilon_{k+1}^z|^2 - 2(\varepsilon_{k+1}^z)^\top r_{k+1}^i) + \sum_{i=1}^N (|r_{k+1}^i + (z_{k+1} - z_k)|^2) + 2 \sum_{i=1}^N (\varepsilon_{k+1}^z - \varepsilon_k^z)^\top (x_{k+1}^i - x_*^i) - 2 \sum_{i=1}^N (\varepsilon_{k+1}^{d,i} + \varepsilon_{k+1}^z)^\top (z_{k+1} - z_*).$$

Modifying the error terms using $x_{k+1}^i - x_*^i = r_{k+1}^i + (z_{k+1} - z_*)$ and $r_{k+1}^i = u_{k+1}^i - u_k^i + \varepsilon_{k+1}^z$, we get,

$$\begin{aligned} N(V_{k+1} - V_k) \leq & - \sum_{i=1}^N |r_{k+1}^i + (z_{k+1} - z_k)|^2 + \sum_{i=1}^N \left(-2(\varepsilon_{k+1}^z - \varepsilon_k^z)^\top (u_{k+1}^i - u_*^i) - |\varepsilon_{k+1}^z|^2 \right. \\ & \left. + 2(\varepsilon_{k+1}^{d,i} + \varepsilon_k^z)^\top (z_{k+1} - z_*) - 2(\varepsilon_k^z)^\top (u_k^i - u_*^i) - 2(\varepsilon_k^z)^\top \varepsilon_{k+1}^z \right). \end{aligned} \quad (21)$$

Furthermore, we expand the squared norm term:

$$- \sum_{i=1}^N |r_{k+1}^i + (z_{k+1} - z_k)|^2 = - \sum_{i=1}^N |r_{k+1}^i|^2 - 2 \sum_{i=1}^N (r_{k+1}^i)^\top (z_{k+1} - z_k) - \sum_{i=1}^N |z_{k+1} - z_k|^2. \quad (22)$$

Next we will establish a bound for the cross term $-2 \sum_{i=1}^N (r_{k+1}^i)^\top (z_{k+1} - z_k)$.

We recall that (13) implies (16). This minimization property leads to the following pair of inequalities:

$$\begin{aligned} g(z_{k+1}) - g(z_k) & \leq \rho \sum_{i=1}^N (u_{k+1}^i + \varepsilon_{k+1}^{d,i} + \varepsilon_{k+1}^z)^\top (z_{k+1} - z_k) \\ g(z_k) - g(z_{k+1}) & \leq -\rho \sum_{i=1}^N (u_k^i + \varepsilon_k^{d,i} + \varepsilon_k^z)^\top (z_{k+1} - z_k). \end{aligned}$$

By adding these inequalities and rearranging terms, we obtain,

$$0 \leq \sum_{i=1}^N (u_{k+1}^i - u_k^i + \varepsilon_{k+1}^{d,i} - \varepsilon_k^{d,i} + \varepsilon_{k+1}^z - \varepsilon_k^z)^\top (z_{k+1} - z_k) = \sum_{i=1}^N (r_{k+1}^i + \varepsilon_{k+1}^{d,i} - \varepsilon_k^{d,i} - \varepsilon_k^z)^\top (z_{k+1} - z_k).$$

This yields the desired bound on the cross term $-2 \sum_{i=1}^N (r_{k+1}^i)^\top (z_{k+1} - z_k)$ which takes the form

$$-2 \sum_{i=1}^N (r_{k+1}^i)^\top (z_{k+1} - z_k) \leq 2 \sum_{i=1}^N (\varepsilon_{k+1}^{d,i} - \varepsilon_k^{d,i} - \varepsilon_k^z)^\top (z_{k+1} - z_k). \quad (23)$$

As the next step, we rewrite $|z_{k+1} - z_k|$ in terms of the gradients of f^i and g . This will be important for deriving the desired convergence result. From the first-order optimality condition of the z -update (13), and rearranging for z_{k+1} , we get,

$$z_{k+1} \in \frac{1}{N} \sum_{i=1}^N \left(x_{k+1}^i + u_k^i + \varepsilon_{k+1}^{d,i} \right) - \frac{1}{\rho N} \partial g(z_{k+1}). \quad (24)$$

The optimality condition for x_{k+1}^i (see (12)) results in,

$$-z_k = \varepsilon_k^z + \frac{1}{N} \sum_{i=1}^N \left(-x_{k+1}^i - u_k^i - \frac{1}{\rho} \nabla f^i(x_{k+1}^i) \right). \quad (25)$$

We combine (24) and (25) to express the update for $z_{k+1} - z_k$ as follows,

$$z_{k+1} - z_k \in -\frac{1}{\rho N} \left(\sum_{i=1}^N \nabla f^i(x_{k+1}^i) + \partial g(z_{k+1}) \right) + \varepsilon_k^z + \frac{1}{N} \sum_{i=1}^N \varepsilon_{k+1}^{d,i}.$$

Thus, $z_{k+1} - z_k$ depends on the averaged gradients ∇f^i and ∂g , scaled by the penalty parameter ρ . We now take the square and apply Young's inequality on the cross term with $\gamma' = 2$, which yields,

$$-|z_{k+1} - z_k|^2 \leq -\frac{1}{2} \left| \frac{1}{\rho N} \left(\sum_{i=1}^N \nabla f^i(x_{k+1}^i) + \nu_{k+1} \right) \right|^2 + \left| \varepsilon_k^z + \frac{1}{N} \sum_{i=1}^N \varepsilon_{k+1}^{d,i} \right|^2, \quad (26)$$

for any $\nu_{k+1} \in \partial g(z_{k+1})$.

By substituting (22), (23), and (26) in (21), and applying Young's inequality to cross terms, we derive the following inequality:

$$\begin{aligned} N(V_{k+1} - V_k) &\leq - \sum_{i=1}^N |r_{k+1}^i|^2 - \frac{1}{2} \left| \frac{1}{\rho N} \left(\sum_{i=1}^N \nabla f^i(x_{k+1}^i) + \nu_{k+1} \right) \right|^2 + \left| \varepsilon_k^z + \frac{1}{N} \sum_{i=1}^N \varepsilon_{k+1}^{d,i} \right|^2 \\ &\quad + \sum_{i=1}^N \left(\gamma_4 |\varepsilon_{k+1}^z - \varepsilon_k^z|^2 + \frac{1}{\gamma_4} |r_{k+1}^i|^2 + \gamma_1 |\varepsilon_{k+1}^z|^2 + \frac{1}{\gamma_1} |u^k - u^*|^2 - |\varepsilon_{k+1}^z|^2 + \gamma_2 |2\varepsilon_{k+1}^{d,i} - \varepsilon_k^{d,i}|^2 \right. \\ &\quad \left. + \frac{1}{\gamma_2} |z_{k+1} - z_k|^2 + \gamma_3 |\varepsilon_{k+1}^{d,i} + \varepsilon_k^z|^2 + \frac{1}{\gamma_3} |z_k - z_*|^2 + 2(\varepsilon_{k+1}^z - 2\varepsilon_k^z)^\top \varepsilon_{k+1}^z \right), \end{aligned}$$

for any $\nu_{k+1} \in \partial g(z_{k+1})$.

We can further simplify the expression by choosing the values as $\gamma_1 = \gamma_3 = \gamma_k$ and $\gamma_2 = \gamma_4 = 3$,

$$\begin{aligned} N(V_{k+1} - V_k) &\leq \frac{N}{\gamma_k} V_k - \frac{2}{3} \sum_{i=1}^N |r_{k+1}^i|^2 - \frac{1}{6} \left| \frac{1}{\rho N} \left(\sum_{i=1}^N \nabla f^i(x_{k+1}^i) + \nu_{k+1} \right) \right|^2 + \left| \varepsilon_k^z + \frac{1}{N} \sum_{i=1}^N \varepsilon_{k+1}^{d,i} \right|^2 \\ &\quad + \sum_{i=1}^N \left(3|\varepsilon_{k+1}^z - \varepsilon_k^z|^2 + \gamma_k |\varepsilon_{k+1}^z|^2 - |\varepsilon_{k+1}^z|^2 + 3|2\varepsilon_{k+1}^{d,i} - \varepsilon_k^{d,i}|^2 + \gamma_k |\varepsilon_{k+1}^{d,i} + \varepsilon_k^z|^2 \right. \\ &\quad \left. + 4|\varepsilon_{k+1}^z - 2\varepsilon_k^z|^2 + 4|\varepsilon_{k+1}^z|^2 \right), \end{aligned}$$

for any $\nu_{k+1} \in \partial g(z_{k+1})$.

The error values arising from event-based communication are bounded by the communication thresholds, $|\varepsilon_k^{d,i}| \leq \Delta_k^d$, $|\varepsilon_k^z| \leq \Delta_k^z$. This leads to the following inequality,

$$\begin{aligned} V_{k+1} &\leq \left(1 + \frac{1}{\gamma_k} \right) V_k - \frac{2}{3N} \sum_{i=1}^N |r_{k+1}^i|^2 - \frac{1}{6N} \left| \frac{1}{\rho N} \left(\sum_{i=1}^N \nabla f^i(x_{k+1}^i) + \nu_{k+1} \right) \right|^2 \\ &\quad + (3\gamma_k + 51 + \frac{8}{3N}) \Delta_k^z{}^2 + (2\gamma_k + 30 + \frac{8}{3N}) \Delta_k^d{}^2, \end{aligned}$$

for any $\nu_{k+1} \in \partial g(z_{k+1})$, where $r_{k+1}^i = x_{k+1}^i - z_{k+1}$ represents the residuals at step $k+1$. Simplifying the expression, we obtain:

$$V_{k+1} \leq \left(1 + \frac{1}{\gamma_k} \right) V_k - \frac{2}{3N} \sum_{i=1}^N |x_{k+1}^i - z_{k+1}|^2 - \frac{1}{6N} \left| \frac{1}{\rho N} \left(\sum_{i=1}^N \nabla f^i(x_{k+1}^i) + \nu_{k+1} \right) \right|^2 + \mathcal{O}(\gamma_k \cdot \Delta_k^2), \quad (27)$$

for any $\nu_{k+1} \in \partial g(z_{k+1})$, where Δ_k represents the time-varying communication threshold, i.e., chosen bound for the perturbation term. This inequality suggests a relationship between V_k and V_{k+1} at consecutive steps.

To analyze the convergence of the sequence V_k , we apply Polyak's Lemma (Lemma 2 in (Polyak, 1987, Chapter 2.2)), which establishes convergence under certain additional assumptions. Polyak's Lemma states that if a sequence V_k satisfies an inequality of the form,

$$V_{k+1} \leq \left(1 + \frac{1}{\gamma_k} \right) V_k - c_k^- + c_k^+,$$

where c_k^- and c_k^+ are sequences of non-negative terms, then the sequence V_k is bounded above provided that $\sum_{k=0}^{\infty} \frac{1}{\gamma_k} < \infty$, and $\sum_{k=0}^{\infty} c_k^+ < \infty$.

We choose $\gamma_k = (k+1)^p$ for some $p > 1$ to ensure convergence. Substituting this choice into the recurrence relation, we obtain:

$$V_{k+1} \leq \left(1 + \frac{1}{(k+1)^p}\right) V_k - \frac{2}{3N} \sum_{i=1}^N |x_{k+1}^i - z_{k+1}|^2 - \frac{1}{6N} \left| \frac{1}{\rho N} \left(\sum_{i=1}^N \nabla f^i(x_{k+1}^i) + \nu_{k+1} \right) \right|^2 + \mathcal{O}(k^p \cdot \Delta_k^2),$$

for any $\nu_{k+1} \in \partial g(z_{k+1})$.

By assumption the communication threshold Δ_k decays as $\Delta_k \sim \mathcal{O}(1/k^t)$. Under this assumption, the perturbation term $\mathcal{O}(k^p \cdot \Delta_k^2)$ scales as $\mathcal{O}(k^{p-2t} \cdot \Delta_0^2)$, which satisfies $\sum_{k=0}^{\infty} \gamma_k \Delta_k^2 < \infty$ if $p > 1$ and $t > \frac{1+p}{2}$. These conditions ensure that the perturbation term decays sufficiently fast, ensuring boundedness of V_k .

Finally, by summing over $k = 0$ to K and dividing by $K+1$, we obtain the following bound for the average of the residuals and gradient terms:

$$\frac{1}{K+1} \sum_{k=0}^K \left(\frac{2}{3N} \sum_{i=1}^N |x_{k+1}^i - z_{k+1}|^2 + \frac{1}{6N} \left| \frac{1}{\rho N} \left(\sum_{i=1}^N \nabla f^i(x_{k+1}^i) + \nu_{k+1} \right) \right|^2 \right) \leq \mathcal{O}\left(\frac{1}{K}\right), \quad (28)$$

for any $\nu_{k+1} \in \partial g(z_{k+1})$, where communication threshold decays $\Delta_k \leq \frac{\Delta_0}{(k+1)^2}$. This result establishes a sublinear convergence rate for both the residuals and the gradient terms. The rate of decay of the communication threshold ensures that these error terms decrease at a rate proportional to $\mathcal{O}\left(\frac{1}{K}\right)$, which yields the desired result. \square

C. Derivation of Alg. 2 as a Dynamical System

In this section, we represent Alg. 2 as a dynamical system that consists of linear dynamics with a nonlinear feedback interconnection. The communication structure is summarized with Fig. 2.

For the convenience of the reader, we start by restating Alg. 2, which is based on an over-relaxed ADMM algorithm.

Algorithm 3 Event-Based Distributed Optimization with Over-Relaxed ADMM

Require: Functions f and g , matrices A and B , vector c , parameters ρ and α

Require: Initial condition x_0, z_0

$$r_0 = \hat{r}_0^s = \hat{r}_0^u = Ax_0, \quad s_0 = \hat{s}_0^r = \hat{s}_0^u = Bz_0, \quad u_0 = \hat{u}_0^r = \hat{u}_0^s = 0$$

for $k = 0$ to t_{\max} **do**

$$\begin{aligned} \hat{s}_k^r, \hat{u}_k^r &\leftarrow \text{event-based receive of } s_{k+1} - s_{[k]}, u_{k+1} - u_{[k]} \\ x_{k+1} &= \arg \min_x f(x) + \frac{\rho}{2} |Ax + \hat{s}_k^r - c + \hat{u}_k^r|^2 && \{\text{r-agent}\} \\ \text{event-based send of } r_{k+1} - r_{[k]} &\text{ where } r_{k+1} = Ax_{k+1} \end{aligned}$$

$$\begin{aligned} \hat{r}_{k+1}^s, \hat{u}_k^s &\leftarrow \text{event-based receive of } r_{k+1} - r_{[k]}, u_{k+1} - u_{[k]} \\ z_{k+1} &= \arg \min_z g(z) + \frac{\rho}{2} |\alpha \hat{r}_{k+1}^s - (1 - \alpha)Bz_k + Bz - \alpha c + \hat{u}_k^s|^2 && \{\text{s-agent}\} \\ \text{event-based send of } s_{k+1} - s_{[k]} &\text{ where } s_{k+1} = Bz_{k+1} \end{aligned}$$

$$\begin{aligned} \hat{r}_{k+1}^u, \hat{s}_{k+1}^u &\leftarrow \text{event-based receive of } r_{k+1} - r_{[k]}, s_{k+1} - s_{[k]} \\ u_{k+1} &= u_k + \alpha \hat{r}_{k+1}^u - (1 - \alpha) \hat{s}_k^u + \hat{s}_{k+1}^u - \alpha c && \{\text{u-agent}\} \\ \text{event-based send of } u_{k+1} - u_{[k]} & \end{aligned}$$

if $\text{mod}(k + 1, T) = 0$ **then**

$$\text{reset} \rightarrow \hat{r}_{k+1}^{u;s} = r_{k+1}, \hat{s}_{k+1}^{u;r} = s_{k+1}, u_{k+1}^{r;s} = u_{k+1}$$

end if

end for

The following definitions will be useful for simplifying the updates of the iterates:

Definition C.1. Let Assumption 3.1 and 3.2 hold. We define the function $\hat{f} : \mathbb{R}^n \rightarrow \mathbb{R}$ as follows,

$$\hat{f} = (\rho^{-1}f) \circ A^{-1}, \quad (29)$$

where ρ is the step-size of Alg. 2. The function is $\hat{m} := m/(\rho\bar{\sigma}^2(A))$ -strongly convex and $\hat{L} := L/(\rho\sigma^2(A))$ -smooth, and has therefore the condition number

$$\kappa := \frac{\hat{L}}{\hat{m}} = \frac{L}{m} \frac{\bar{\sigma}^2(A)}{\sigma^2(A)}.$$

Definition C.2. Let Assumption 3.1 and 3.2 hold. The function $\hat{g} : \mathbb{R}^m \rightarrow \bar{\mathbb{R}}$ is defined as

$$\hat{g} = (\rho^{-1}g) \circ B^\dagger + \psi_{\text{im}(B)}, \quad (30)$$

where B^\dagger is the Moore-Penrose inverse of B , $\psi_{\text{im}(B)}$ is the indicator function of the image of B , and ρ is the step-size of Alg. 2.

We proceed by summarizing the notation that will be used subsequently. The sequences r_k and s_k are defined as $r_k := Ax_k$ and $s_k := Bz_k$. We introduced the variable \hat{r}_k^s , for example, which models agent s 's estimate of the variable r_k . The variables $\hat{r}_k^u, \hat{s}_k^r, \hat{s}_k^u$, etc., are defined analogously and follow the notational convention

$$\widehat{\text{variable}}_k^{\text{receiving_agent}} := \text{receiving_agent's estimate of variable at time } k.$$

As a result of the event-based communication, the local estimates $\hat{r}_k^s, \hat{r}_k^u, \hat{s}_k^r$, etc., differ from r_k, s_k , etc. These differences will be captured by the variable ε for which we introduce the following notational convention:

$$\begin{aligned} \varepsilon_k^{\text{variable, receiving_agent}} &:= \widehat{\text{variable}}_k^{\text{receiving_agent}} - \text{variable}_k \\ &= \text{receiving_agent's estimation error of variable at time } k. \end{aligned}$$

We further introduce the error

$$e_k := (\varepsilon_{k+1}^{rs}, \varepsilon_{k+1}^{ru}, \varepsilon_{k+1}^{su}, \varepsilon_k^{sr}, \varepsilon_k^{su}, \varepsilon_k^{ur}, \varepsilon_k^{us}), \quad (31)$$

that collects the estimation errors of the different agents. By virtue of the event-based communication mechanism and the reset mechanism, the error e_k is bounded by the communication threshold Δ . We finally introduce the notation for the corresponding communication thresholds, Δ^{rs} , Δ^{sr} , etc. (see Fig. 2), according to the same rationale:

$$\Delta^{\text{variable, receiving_agent}} := \text{threshold for triggering a communication of variable to receiving_agent.}$$

To sum up, the vector e_k contains the errors on the communication lines shown on Fig. 2. For example, ε_{k+1}^{rs} stands for the difference between the actual value of state r_{k+1} and agent s 's estimate \hat{r}_{k+1}^s , that is, $\varepsilon_{k+1}^{rs} = \hat{r}_{k+1}^s - r_{k+1}$ at time step $k+1$.

If the value r_{k+1} has deviated more than Δ^{rs} amount since the time-step $[k]$, where the last value $r_{[k]}$ has been communicated to the agent s , a communication is triggered. This means $rs \in \mathcal{C}_{k+1}$.

$$|r_{k+1} - r_{[k]}^{rs}| > \Delta^{rs} \iff rs \in \mathcal{C}_{k+1} \iff [k+1] = k+1.$$

The set \mathcal{D}_{k+1} , which is a subset of \mathcal{C}_{k+1} , collects indices of failed transmission lines at time step $k+1$. We further introduce that the superscript c to denote the complement of a set. If the communication does not fail, that is, $rs \in \mathcal{D}_{k+1}^c$, then agent s 's estimate of r_{k+1} is updated as follows.

$$rs \in \mathcal{C}_{k+1} \wedge rs \in \mathcal{D}_{k+1}^c \iff \hat{r}_{k+1}^s = \hat{r}_k^s + (r_{k+1} - r_{[k]}).$$

To incorporate the effect of communication drops, we introduce the variable χ_{k+1}^{rs} , which represents the disturbance that results from dropped communications,

$$rs \in \mathcal{D}_{k+1} \Rightarrow \chi_{k+1}^{rs} = -(r_{k+1} - r_{[k]}). \quad (32)$$

Therefore, the dynamics of \hat{r}_{k+1}^s are expressed as follows

$$\hat{r}_{k+1}^s = r_{[k+1]} + \sum_{l=1}^{k+1} \chi_l^{rs}. \quad (33)$$

When deriving the previous equation, we have exploited the fact that

$$\begin{aligned} rs \in \mathcal{C}_{k+1} &\Rightarrow [k+1] = k+1 \\ rs \in \mathcal{C}_{k+1}^c &\Rightarrow [k+1] = [k]. \end{aligned}$$

To summarize, in the case of communication drop, the agent s updates the image of r with a disturbed value.

We now express the different minimization steps in Alg. 3 by their corresponding stationarity conditions, and simplify the corresponding expressions. The minimization step for the primal variable x can be rewritten as follows

$$\begin{aligned} x_{k+1} &= \arg \min_{x \in \mathbb{R}^p} f(x) + \frac{\rho}{2} |Ax + \hat{s}_k^r - c + \hat{u}_k^r|^2 \\ &= A^{-1} \arg \min_{r \in \mathbb{R}^n} f(A^{-1}r) + \frac{\rho}{2} |r + \hat{s}_k^r - c + \hat{u}_k^r|^2, \end{aligned}$$

due to the fact that A is invertible, which yields

$$r_{k+1} = \arg \min_{r \in \mathbb{R}^n} \hat{f}(r) + \frac{1}{2} |r + \hat{s}_k^r - c + \hat{u}_k^r|^2.$$

The variable r_{k+1} satisfies therefore the following stationarity condition

$$0 = \nabla \hat{f}(r_{k+1}) + r_{k+1} + \hat{s}_k^r - c + \hat{u}_k^r,$$

which can be rearranged to

$$r_{k+1} - c = -\nabla \hat{f}(r_{k+1}) - s_k - \varepsilon_k^{sr} - u_k - \varepsilon_k^{ur}, \quad (34)$$

where \hat{s}_k^r is replaced by $s_k + \varepsilon_k^{sr}$ and \hat{u}_k^r by $u_k + \varepsilon_k^{ur}$.

Similarly, the update step of the auxiliary variable z can be reformulated as

$$\begin{aligned} z_{k+1} &= \arg \min_{z \in \mathbb{R}^n} g(z) + \frac{\rho}{2} |\alpha \hat{r}_{k+1}^s - (1 - \alpha)s_k + Bz - \alpha c + \hat{u}_k^s|^2 \\ &= B^\dagger \arg \min_{s \in \mathbb{R}^m} g(B^\dagger s) + \psi_{\text{im}(B)}(s) + \frac{\rho}{2} |\alpha \hat{r}_{k+1}^s - (1 - \alpha)s_k + s - \alpha c + \hat{u}_k^s|^2, \end{aligned}$$

since the matrix B has full column rank and therefore possesses the left inverse B^\dagger . This yields

$$s_{k+1} = \arg \min_{s \in \mathbb{R}^m} \hat{g}(s) + \frac{1}{2} |\alpha \hat{r}_{k+1}^s - (1 - \alpha)s_k + s - \alpha c + \hat{u}_k^s|^2,$$

and implies the following stationarity condition for s_{k+1}

$$0 \in \partial \hat{g}(s_{k+1}) + \alpha \hat{r}_{k+1}^s - (1 - \alpha)s_k + s_{k+1} - \alpha c + \hat{u}_k^s. \quad (35)$$

This stationarity condition can be reformulated as

$$s_{k+1} = s_k + (\alpha - 1)u_k + \alpha \nabla \hat{f}(r_{k+1}) - \gamma_{k+1} - \varepsilon_k^{us} + \alpha \varepsilon_k^{sr} + \alpha \varepsilon_k^{ur} - \alpha \varepsilon_{k+1}^{rs}, \quad (36)$$

for some $\gamma_{k+1} \in \partial \hat{g}(s_{k+1})$, and where we have expressed \hat{r}_{k+1}^s as $r_{k+1} + \varepsilon_{k+1}^{rs}$ and \hat{u}_k^s as $u_k + \varepsilon_k^{us}$. We have further replaced $r_{k+1} - c$ by the expression given in (34).

The update of the dual variables u_k evolve according to the following dynamics:

$$\begin{aligned} u_{k+1} &= u_k + \alpha \hat{r}_{k+1}^u - (1 - \alpha)\hat{s}_k^u + \hat{s}_{k+1}^u - \alpha c \\ &= u_k + \alpha(r_{k+1} + \varepsilon_{k+1}^{ru}) - (1 - \alpha)(s_k + \varepsilon_k^{su}) + (s_{k+1} + \varepsilon_{k+1}^{su}) - \alpha c. \end{aligned}$$

The dynamics can be further simplified by replacing s_{k+1} with the help of (35), which yields:

$$u_{k+1} = -\gamma_{k+1} - \alpha \varepsilon_{k+1}^{rs} + \alpha \varepsilon_{k+1}^{ru} + \varepsilon_{k+1}^{su} + (\alpha - 1)\varepsilon_k^{su} - \varepsilon_k^{us}. \quad (37)$$

As a result of these simplifications, we note that r_{k+1} is uniquely determined by s_k , u_k and the corresponding errors ε_k^{sr} and ε_k^{ur} . We further note that according to (36) and (37) the iterates of Alg. 3 can be represented as an interconnection between a linear dynamical system, with a nonlinear feedback interconnection that models the evaluation of the gradient $\nabla \hat{f}$ and $\partial \hat{g}$. The state of the dynamical system is therefore chosen as $\xi_k := (s_k, u_k)$, the output as $y_k := (r_{k+1} - c, s_{k+1})$, and the input as $v_k := (\nabla \hat{f}(r_{k+1}), \gamma_{k+1})$, where $\gamma_{k+1} \in \partial \hat{g}(s_{k+1})$. We also define output variables $w_k^1 := (r_{k+1} - c, \nabla \hat{f}(r_{k+1}))$, $w_k^2 := (s_{k+1}, \gamma_{k+1})$, which will be employed for the convergence analysis.

According to these definitions, we can express the iterates of Alg. 3 as trajectories of the following nonlinear dynamical system,

$$\xi_{k+1} = \underbrace{\begin{bmatrix} 1 & \alpha - 1 \\ 0 & 0 \end{bmatrix}}_{:=\hat{A}} \xi_k + \underbrace{\begin{bmatrix} \alpha & -1 \\ 0 & -1 \end{bmatrix}}_{:=\hat{B}} v_k + \underbrace{\begin{bmatrix} -\alpha & 0 & 0 & \alpha & 0 & \alpha & -1 \\ -\alpha & \alpha & 1 & 0 & \alpha & -1 & 0 & -1 \end{bmatrix}}_{:=\hat{E}} e_k, \quad v_k = \phi(y_k), \quad (38)$$

$$y_k = \underbrace{\begin{bmatrix} -1 & -1 \\ 1 & \alpha - 1 \end{bmatrix}}_{:=\hat{C}} \xi_k + \underbrace{\begin{bmatrix} -1 & 0 \\ \alpha & -1 \end{bmatrix}}_{:=\hat{D}} v_k + \underbrace{\begin{bmatrix} 0 & 0 & 0 & -1 & 0 & -1 & 0 \\ -\alpha & 0 & 0 & \alpha & 0 & \alpha & -1 \end{bmatrix}}_{:=\hat{E}^y} e_k,$$

$$w_k^1 = \underbrace{\begin{bmatrix} -1 & -1 \\ 0 & 0 \end{bmatrix}}_{:=\hat{C}^1} \xi_k + \underbrace{\begin{bmatrix} -1 & 0 \\ 1 & 0 \end{bmatrix}}_{:=\hat{D}^1} v_k + \underbrace{\begin{bmatrix} 0 & 0 & 0 & -1 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}}_{:=\hat{E}^1} e_k,$$

$$w_k^2 = \underbrace{\begin{bmatrix} 1 & \alpha - 1 \\ 0 & 0 \end{bmatrix}}_{:=\hat{C}^2} \xi_k + \underbrace{\begin{bmatrix} \alpha & -1 \\ 0 & 1 \end{bmatrix}}_{:=\hat{D}^2} v_k + \underbrace{\begin{bmatrix} -\alpha & 0 & 0 & \alpha & 0 & \alpha & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}}_{:=\hat{E}^2} e_k, \quad (39)$$

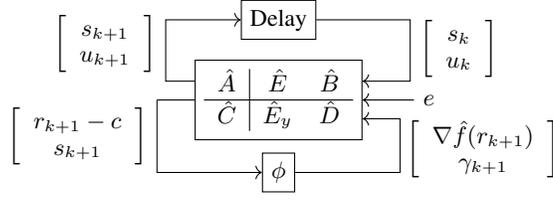


Figure 7: The dynamical system following (38) is visualized.

where ϕ denotes the nonlinear feedback interconnection that captures the evaluation of the gradients $\nabla \hat{f}$ and $\partial \hat{g}$. Fig. 7 provides a graphical representation of the time-invariant dynamics determined by the matrices $\hat{A}, \hat{B}, \hat{C}, \hat{D}$.

We close the section with the following proposition that shows that $|e_k|$ is bounded.

Proposition C.3. *The error e_k at iteration k is bounded by*

$$|e_k| \leq \Delta, \quad \Delta := \sum_{l \in \{rs, ru, su, sr, su, ur, us\}} \Delta^l + T \bar{\chi}^l,$$

where the variable $\bar{\chi}^l$ is an upper bound on the communication drops.

Proof. The proof is analogous to Prop. 2.1. The error resulting from the event-based communication structure is given by

$$e_{k+1}^{rs} = \hat{r}_{k+1}^s - r_{k+1} = \underbrace{r_{[k+1]} - r_{k+1}}_{\text{I}} + \underbrace{\sum_{l=1}^{k+1} \chi_l^{rs}}_{\text{II}}. \quad (40)$$

We further note that the first term is bounded by Δ^{rs} by virtue of the communication rule

$$|r_{k+1} - r_{[k+1]}| \leq \Delta^{rs}.$$

Through the assumption $|\chi_l^{rs}| \leq \bar{\chi}^{rs}$, the second part is bounded by $T \bar{\chi}^{rs}$, where T is the reset period. Therefore, we conclude that $|e_{k+1}^{rs}|$ is bounded by $\Delta^{rs} + T \bar{\chi}^{rs}$. Similarly, the other elements of the vector e_k are bounded by $\Delta^{ru} + T \bar{\chi}^{ru}$, $\Delta^{su} + T \bar{\chi}^{su}$, etc. Hence, $|e_k|$ is bounded by Δ where

$$\Delta = \sum_{l \in \{rs, ru, su, sr, su, ur, us\}} \Delta^l + T \bar{\chi}^l.$$

□

The analysis indicates that a periodic reset with a period T is required to achieve a bounded error. If no reset is included, Alg. 2 may not converge, which could result in a large error that accumulates over time. The dependence of Δ on the period T highlights how the reset period T affects the error (where smaller T leads to a smaller error bound). If there are no communication failures, there is also no need for a reset ($\bar{\chi} = 0$), and Δ reduces to the collection of communication thresholds.

D. Convergence Analysis

We first start by proving the following intermediate lemmas.

Lemma D.1. *Let Assumption 3.2 be satisfied. Then, the following holds,*

$$[(r_1 - r_2)^\top (\nabla \hat{f}(r_1) - \nabla \hat{f}(r_2))^\top] \left(\begin{bmatrix} -2\hat{m}\hat{L} & (\hat{m} + \hat{L}) \\ (\hat{m} + \hat{L}) & -2 \end{bmatrix} \otimes I_n \right) \begin{bmatrix} r_1 - r_2 \\ \nabla \hat{f}(r_1) - \nabla \hat{f}(r_2) \end{bmatrix} \geq 0, \quad (41)$$

for all $r_1, r_2 \in \mathbb{R}^n$.

Proof. We define the auxiliary function $\tilde{f}(r) := \hat{f}(r) - \frac{\hat{m}}{2}|r|^2$, which is $\hat{L} - \hat{m}$ -smooth and convex by the properties of \hat{f} . Then the following inequality holds,

$$(\nabla \tilde{f}(r_1) - \nabla \tilde{f}(r_2))^\top (r_1 - r_2) \geq \frac{1}{\hat{L} - \hat{m}} |\nabla \tilde{f}(r_1) - \nabla \tilde{f}(r_2)|^2,$$

for any $r_1, r_2 \in \mathbb{R}^n$. Substituting $\tilde{f}(r) := \hat{f}(r) - \frac{\hat{m}}{2}|r|^2$ and $\nabla \tilde{f}(r) = \nabla \hat{f}(r) - \hat{m}r$, we get

$$(\hat{m} + \hat{L})(r_1 - r_2)^\top (\nabla \hat{f}(r_1) - \nabla \hat{f}(r_2)) \geq \hat{m}\hat{L}|r_1 - r_2|^2 + |\nabla \hat{f}(r_1) - \nabla \hat{f}(r_2)|^2,$$

which yields the desired result. \square

Lemma D.2. *Let Assumption 3.2 be satisfied. Then, the following holds,*

$$[(s_1 - s_2)^\top (\gamma_1 - \gamma_2)^\top] \left(\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \otimes I_m \right) \begin{bmatrix} s_1 - s_2 \\ \gamma_1 - \gamma_2 \end{bmatrix} \geq 0, \quad (42)$$

where $\gamma_1 \in \partial \hat{g}(s_1)$ and $\gamma_2 \in \partial \hat{g}(s_2)$ and for any $s_1, s_2 \in \mathbb{R}^m$.

Proof. The subdifferential of a convex function is a monotone operator, and therefore

$$(s_1 - s_2)^\top (\gamma_1 - \gamma_2) \geq 0.$$

\square

Lemma D.3. *Let x_*, z_* denote the minimizer of (3) and define $r_* := Ax_*$, $s_* := Bz_*$, $\beta_* := \nabla \hat{f}(r_*)$, and $\gamma_* \in \partial \hat{g}(s_*)$. Then, the iterates of Alg. 3 with step-size $\rho = \rho_0(\hat{m}\hat{L})^{\frac{1}{2}}$ satisfy*

$$(w_k^i - w_*^i)^\top M^i (w_k^i - w_*^i) \geq 0, \quad \forall i \in \{1, 2\}, \quad \forall k \geq 0,$$

with

$$M^1 := \begin{bmatrix} -2\rho_0^{-2} & \rho_0^{-1}(\kappa^{-\frac{1}{2}} + \kappa^{\frac{1}{2}}) \\ \rho_0^{-1}(\kappa^{-\frac{1}{2}} + \kappa^{\frac{1}{2}}) & -2 \end{bmatrix} \otimes I_n, \quad M^2 := \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \otimes I_m,$$

where

$$w_k^1 := \begin{bmatrix} r_{k+1} - c \\ \beta_{k+1} \end{bmatrix}, \quad w_k^2 := \begin{bmatrix} s_{k+1} \\ \gamma_{k+1} \end{bmatrix}, \quad w_*^1 := \begin{bmatrix} r_* - c \\ \beta_* \end{bmatrix}, \quad w_*^2 := \begin{bmatrix} s_* \\ \gamma_* \end{bmatrix}.$$

Proof. The proof follows directly from Lemma D.1 and Lemma D.2. \square

Lemma D.4. *Let the sequence $V_k \geq 0$ satisfy*

$$V_{k+1} \leq V_k(1 - \tilde{\alpha}) + \tilde{\beta}\tilde{\alpha}, \quad (43)$$

for all $k \geq 0$, where the parameters $\tilde{\alpha}, \tilde{\beta}$ satisfy $0 < \tilde{\alpha} < 1$ and $0 \leq \tilde{\beta}$. Then, the following holds for all $k \geq 0$:

$$V_k \leq V_0(1 - \tilde{\alpha})^k + \tilde{\beta}. \quad (44)$$

Proof. We prove the lemma by induction.

The claim holds for $k = 0$. We therefore assume that the claim holds for k and show that, as a result, the claim holds for $k + 1$. More precisely,

$$\begin{aligned} V_{k+1} &\leq V_k(1 - \tilde{\alpha}) + \tilde{\beta}\tilde{\alpha} \\ &\leq V_0(1 - \tilde{\alpha})^{k+1} + (1 - \tilde{\alpha})\tilde{\beta} + \tilde{\beta}\tilde{\alpha} \\ &\leq V_0(1 - \tilde{\alpha})^{k+1} + \tilde{\beta}, \end{aligned} \quad (45)$$

which completes the induction argument. \square

In App. C, we expressed the iterates of Alg. 3 as the trajectories of a dynamical system. The dynamical system was given as a linear time-invariant system that was interconnected in feedback with a nonlinear function ϕ . We now arrive at the main result that will be used to show convergence of Alg. 3.

Theorem D.5. *Let Assumption 3.2 be satisfied, let the step-size for Alg. 3 be $\rho = \rho_0(\hat{m}\hat{L})^{\frac{1}{2}}$, and let $\xi_* = (Bz_*, u_*)$, where (x_*, z_*) is the minimizer of (3) and u_* the corresponding dual variable.*

Suppose there exists a positive definite matrix $P \succ 0$, $0 < \tau < 1$, and nonnegative constants $\lambda^1, \lambda^2, \gamma^1, \gamma^2, \gamma^3$ and γ^4 such that the following linear matrix inequality

$$0 \succeq \begin{bmatrix} (1+\gamma^1)\hat{A}^\top P \hat{A} - \tau^2 P & \hat{A}^\top P \hat{B} \\ \hat{B}^\top P \hat{A} & (1+\gamma^2)\hat{B}^\top P \hat{B} \end{bmatrix} + \begin{bmatrix} \hat{C}^1 & \hat{D}^1 \\ \hat{C}^2 & \hat{D}^2 \end{bmatrix}^\top \begin{bmatrix} \Lambda^1 M^1 & 0 \\ 0 & \Lambda^2 M^2 \end{bmatrix} \begin{bmatrix} \hat{C}^1 & \hat{D}^1 \\ \hat{C}^2 & \hat{D}^2 \end{bmatrix} \quad (46)$$

is satisfied, where $\Lambda^1 = \lambda^1(1 + \gamma^3)$, $\Lambda^2 = \lambda^2(1 + \gamma^4)$. Then, for all $k \geq 0$, we have

$$|\xi_k - \xi_*|^2 \leq \kappa_P |\xi_0 - \xi_*|^2 \tau^{2k} + \frac{\bar{\sigma}(Q)\Delta^2}{\underline{\sigma}(P)(1 - \tau^2)}, \quad (47)$$

where $\kappa_P = \bar{\sigma}(P)/\underline{\sigma}(P)$ denotes the condition number of the matrix P , Δ is a bound on the error e_k (see Prop. C.3), and

$$Q = \left(1 + \frac{1}{\gamma^1} + \frac{1}{\gamma^2}\right) \hat{E}^\top P \hat{E} + \left(1 + \frac{1}{\gamma^3} + \frac{1}{\gamma^4}\right) \sum_{i=1}^2 \lambda^i \hat{E}^{i\top} M^i \hat{E}^i. \quad (48)$$

Proof. We consider the following quadratic storage function,

$$V_k = (\xi_k - \xi_*)^\top P (\xi_k - \xi_*),$$

and claim that the following inequality holds for the iterates of Alg. 3:

$$\begin{aligned} V_{k+1} - \tau^2 V_k + \sum_{i=1}^2 \lambda^i (w^i - w_*^i)^\top M^i (w^i - w_*^i) &\leq \\ e_k^\top \left(\left(1 + \frac{1}{\gamma^1} + \frac{1}{\gamma^2}\right) E^\top P E + \sum_{i=1}^2 \lambda^i \left(1 + \frac{1}{\gamma^3} + \frac{1}{\gamma^4}\right) E^{i\top} M^i E^i \right) e_k. \end{aligned}$$

Proof of the claim: We insert the system dynamics stated in App. C into the expression on the left-hand side, which yields

$$\begin{aligned}
 & V_{k+1} - \tau^2 V_k + \sum_{i=1}^2 \lambda^i (w_k^i - w_\star^i)^\top M^i (w_k^i - w_\star^i) \\
 &= (\xi_{k+1} - \xi_\star)^\top P (\xi_{k+1} - \xi_\star) - \tau^2 (\xi_k - \xi_\star)^\top P (\xi_k - \xi_\star) + \sum_{i=1}^2 \lambda^i (w_k^i - w_\star^i)^\top M^i (w_k^i - w_\star^i) \\
 &= \tilde{\xi}_k^\top (A^\top P A - \tau^2 P) \tilde{\xi}_k + \tilde{v}_k^\top \hat{B}^\top P \hat{B} \tilde{v}_k + e_k^\top E^\top P E e_k \\
 &\quad + 2 \left(\tilde{v}_k^\top \hat{B}^\top P \hat{A} \tilde{\xi}_k + e_k^\top E^\top P \hat{A} \tilde{\xi}_k + \tilde{v}_k^\top \hat{B}^\top P E e_k \right) \\
 &\quad + \sum_{i=1}^2 \lambda^i \left(\tilde{\xi}_k^\top \hat{C}^{i\top} M^i \hat{C}^i \tilde{\xi}_k + \tilde{v}_k^\top \hat{D}^{i\top} M^i \hat{D}^i \tilde{v}_k + e_k^\top E^{i\top} M^i E^i e_k \right) \\
 &\quad + 2 \sum_{i=1}^2 \lambda^i \left(\tilde{v}_k^\top \hat{D}^{i\top} M^i \hat{C}^i \tilde{\xi}_k + e_k^\top E^{i\top} M^i \hat{C}^i \tilde{\xi}_k + \tilde{v}_k^\top \hat{D}^{i\top} M^i E^i e_k \right),
 \end{aligned} \tag{49}$$

where $\tilde{\xi}_k = \xi_k - \xi_\star$, and $\tilde{v}_k = v_k - v_\star$, for simplicity. We now apply Young's inequality on the cross terms in (49), which yields

$$\begin{aligned}
 & V_{k+1} - \tau^2 V_k + \sum_{i=1}^2 \lambda^i (w_k^i - w_\star^i)^\top M^i (w_k^i - w_\star^i) \\
 &\leq \tilde{\xi}_k^\top (A^\top P A - \tau^2 P) \tilde{\xi}_k + \tilde{v}_k^\top \hat{B}^\top P \hat{B} \tilde{v}_k + e_k^\top E^\top P E e_k + 2 \tilde{v}_k^\top \hat{B}^\top P \hat{A} \tilde{\xi}_k \\
 &\quad + \gamma^2 (\tilde{v}_k^\top \hat{B}^\top P \hat{B} \tilde{v}_k) + \frac{1}{\gamma^2} (e_k^\top E^\top P E e_k) + \gamma^1 (\tilde{\xi}_k^\top \hat{A}^\top P \hat{A} \tilde{\xi}_k) + \frac{1}{\gamma^1} (e_k^\top E^\top P E e_k) \\
 &\quad + \sum_{i=1}^2 \lambda^i \left(\tilde{\xi}_k^\top \hat{C}^{i\top} M^i \hat{C}^i \tilde{\xi}_k + \tilde{v}_k^\top \hat{D}^{i\top} M^i \hat{D}^i \tilde{v}_k + e_k^\top E^{i\top} M^i E^i e_k + 2 \left(\tilde{v}_k^\top \hat{D}^{i\top} M^i \hat{C}^i \tilde{\xi}_k \right) \right) \\
 &\quad + \sum_{i=1}^2 \lambda^i \left(\gamma^4 \tilde{v}_k^\top \hat{D}^{i\top} M^i \hat{D}^i \tilde{v}_k + \frac{1}{\gamma^4} e_k^\top E^{i\top} M^i E^i e_k + \gamma^3 \tilde{\xi}_k^\top \hat{C}^{i\top} M^i \hat{C}^i \tilde{\xi}_k + \frac{1}{\gamma^3} e_k^\top E^{i\top} M^i E^i e_k \right).
 \end{aligned} \tag{50}$$

If we rearrange the right-hand side of the inequality in matrix form, we obtain,

$$\begin{aligned}
 & V_{k+1} - \tau^2 V_k + \sum_{i=1}^2 \lambda^i (w_k^i - w_\star^i)^\top M^i (w_k^i - w_\star^i) \leq \\
 &\quad \begin{bmatrix} \tilde{\xi}_k^\top & \tilde{v}_k^\top \end{bmatrix} \left(\begin{bmatrix} (1 + \gamma^1) \hat{A}^\top P \hat{A} - \tau^2 P & \hat{A}^\top P \hat{B} \\ \hat{B}^\top P \hat{A} & (1 + \gamma^2) \hat{B}^\top P \hat{B} \end{bmatrix} \right. \\
 &\quad \left. + \begin{bmatrix} \hat{C}^1 & \hat{D}^1 \end{bmatrix}^\top \begin{bmatrix} \lambda^1 M^1 (1 + \gamma^3) & 0 \\ 0 & \lambda^2 M^2 (1 + \gamma^4) \end{bmatrix} \begin{bmatrix} \hat{C}^1 & \hat{D}^1 \\ \hat{C}^2 & \hat{D}^2 \end{bmatrix} \right) \begin{bmatrix} \tilde{\xi}_k \\ \tilde{v}_k \end{bmatrix} \\
 &\quad + e_k^\top \left(\left(1 + \frac{1}{\gamma^1} + \frac{1}{\gamma^2} \right) E^\top P E + \sum_{i=1}^2 \lambda^i \left(1 + \frac{1}{\gamma^3} + \frac{1}{\gamma^4} \right) E^{i\top} M^i E^i \right) e_k.
 \end{aligned} \tag{51}$$

The fact that the linear matrix inequality (46) is satisfied proves the claim. Furthermore, we conclude that $\sum_{i=1}^2 \lambda^i (w_k^i - w_\star^i)^\top M^i (w_k^i - w_\star^i) \geq 0$ from Lemma D.1 and D.2. This simplifies the previous expression to

$$V_{k+1} \leq \tau^2 V_k + e_k^\top Q e_k,$$

where we have also inserted the definition of the matrix Q . The right-hand side can further be bounded by virtue of the reset mechanism and the event-based communication, which results in

$$V_{k+1} \leq \tau^2 V_k + \bar{\sigma}(Q) \Delta^2. \tag{52}$$

We are now in a position where we can apply Lemma D.4, which concludes

$$V_k \leq \tau^{2k} V_0 + \frac{\bar{\sigma}(Q)\Delta^2}{1 - \tau^2}.$$

By definition of the quadratic storage function we conclude

$$|\xi_k - \xi_*|^2 \leq \tau^{2k} \frac{\bar{\sigma}(P)}{\underline{\sigma}(P)} |\xi_0 - \xi_*|^2 + \frac{\bar{\sigma}(Q)\Delta^2}{\underline{\sigma}(P)(1 - \tau^2)},$$

which implies the result of Thm. D.5. \square

We are now ready to prove our main result in Thm. 4.1.

Proof of Thm. 4.1. The proof is based on Thm. D.5 which shows that if the matrix inequality

$$0 \succeq \begin{bmatrix} (1 + \gamma^1)\hat{A}^\top P \hat{A} - \tau^2 P & \hat{A}^\top P \hat{B} \\ \hat{B}^\top P \hat{A} & (1 + \gamma^2)\hat{B}^\top P \hat{B} \end{bmatrix} + \begin{bmatrix} \hat{C}^1 & \hat{D}^1 \\ \hat{C}^2 & \hat{D}^2 \end{bmatrix}^\top \begin{bmatrix} \Lambda^1 M^1 & 0 \\ 0 & \Lambda^2 M^2 \end{bmatrix} \begin{bmatrix} \hat{C}^1 & \hat{D}^1 \\ \hat{C}^2 & \hat{D}^2 \end{bmatrix} \quad (53)$$

is satisfied for a symmetric positive definite matrix P and for positive constants $\Lambda^1, \Lambda^2, \gamma^1, \gamma^2$, the following bound holds

$$|\xi_k - \xi_*|^2 \leq \kappa_P |\xi_0 - \xi_*|^2 \tau^{2k} + \frac{\bar{\sigma}(Q)\Delta^2}{\underline{\sigma}(P)(1 - \tau^2)},$$

where κ_P denotes the condition number of P and Q is defined in App. D. In fact, the following set of parameters satisfies the linear matrix inequality (53),

$$P = \begin{bmatrix} 1 & \alpha - 1 \\ \alpha - 1 & 1 - \frac{1}{\sqrt{\kappa}} \end{bmatrix}, \quad \tau = 1 - \frac{\alpha}{4\kappa^{\epsilon + \frac{1}{2}}}, \quad \Lambda^1 = \alpha \kappa^{\epsilon - \frac{1}{2}}, \quad \Lambda^2 = \alpha, \quad \gamma^1 = \frac{\alpha}{\kappa^{\epsilon + \frac{3}{2}}}, \quad \gamma^2 = \frac{1}{\kappa}.$$

This can be checked as follows: The matrix on the right-hand side of (53) can be expressed as $-\frac{1}{4}\kappa^{-2}\mathbb{L}$, where \mathbb{L} is a symmetric 4×4 matrix (compared to earlier analyses (Nishihara et al., 2015), the last row and last column is not zero). We now prove that \mathbb{L} is positive semidefinite for all sufficiently large κ by checking the leading principle minors, which can be expressed as polynomials in κ . If the leading terms of the principle minors have positive coefficients, it means that for large enough κ , the principle minor will indeed be positive.

The leading term for the first principle minor is given by $6\kappa^{\frac{3}{2} - \epsilon}$ and is therefore positive. Likewise, the second principle minor is dominated by the positive term $24(2 - \alpha)\kappa^{\frac{7}{2} - \epsilon}$. For the third leading principle minor, there are two different cases. If $\epsilon = 0$, the leading term of the third leading principle minor is $16\kappa^5(\alpha^4 - 4\alpha^3 - 4\alpha^2 + 22\alpha - 12)/\alpha$, which is positive for $\alpha \in (0.675, 2)$. If $\epsilon > 0$, the leading term of the third principle minor becomes $192\kappa^5(2 - \alpha)$, which is positive for $\alpha \in (0, 2)$. Finally, for the fourth principle minor, there are also two different cases. If $\epsilon = 0$, the leading term is $64\kappa^{\frac{13}{2}}(\alpha^4 - 4\alpha^3 - 4\alpha^2 + 22\alpha - 12)/\alpha^2$, which is positive for $\alpha \in (0.675, 2)$. If $\epsilon > 0$, the leading term of the fourth principal minor becomes $768\kappa^{\frac{13}{2}}(2 - \alpha)/\alpha$, which is positive for $\alpha \in (0, 2)$. In conclusion, for all sufficiently large κ , all four leading principle minors are positive, which implies that \mathbb{L} is positive definite.

It remains to bound the second term $\bar{\sigma}(Q)/(\underline{\sigma}(P)(1 - \tau^2))$. We again investigate the symbolic expression, and conclude that the term is always bounded by $60\kappa^{2+2\epsilon}/(\alpha(1 - |\alpha - 1|))$ for large enough κ . This concludes the proof. \square

E. Bound on Event-Based Error Variables

We restate Prop. 2.1 and present its proof.

Proposition. *The error $\hat{\zeta}_k - \zeta_k$ at iteration k is bounded by $|\hat{\zeta}_k - \zeta_k| \leq \Delta^d + T\bar{\chi}^d$, where T denotes the reset period (see Alg. 1) and $\bar{\chi}^d$ is a bound on the disturbance χ_k^{di} .*

Proof. We note that the error $\hat{\zeta}_k - \zeta_k$ can be expressed as

$$\hat{\zeta}_k - \zeta_k = \frac{1}{N} \sum_{i=1}^N \underbrace{(d_{[k+1]}^i - d_{k+1}^i)}_I + \underbrace{\sum_{l=T_{[k]}}^k \chi_{l+1}^{di}}_{II},$$

where $T_{[k]}$ denotes the last time instant where a reset has been performed. The terms I and II have each a clear interpretation: In the absence of communication failures, $\hat{\zeta}$ is an average over the primal and dual variables, $x_{[k+1]}^i$ and $u_{[k]}^i$, that were last communicated, which leads to the term I. The term II captures the dropped information through failures. The communication protocol ensures that $|d_{k+1}^i - d_{[k+1]}^i| \leq \Delta^d$, for all $k \geq 0$, which means that the term I is bounded by Δ^d . The bound for the term II arises from the triangle inequality, which yields, $\bar{\chi}^d$ and concludes the proof. \square

The previous proposition required the variable χ_{k+1}^{di} to be bounded. Prop. E.1 establishes such a bound under standard conditions on f and g .

Proposition E.1. *Let f be L -smooth and convex and let $\{z \in \mathbb{R}^n \mid g(z) < \infty\}$ be contained in a ball of radius R . Then, the disturbances χ_k^{di} and χ_k^{zi} are bounded by*

$$|\chi_k^{zi}| \leq 2R, \quad |\chi_k^{di}| \leq (\alpha + 1) \frac{2(\rho + L)}{\rho} |x_*^i| + 2R,$$

for all $i = 1, \dots, N$ and all $k \geq 0$, where $x_*^i := \arg \min_{x \in \mathbb{R}^n} f^i(x) + \rho|x|^2/2$, and where χ_k^{zi} denotes the communication drops when communicating z_k between agent $N + 1$ and agent i .

Proof. Due to the assumption that the domain of g is contained in a ball of radius R , we conclude $|z_k| \leq R$ for all $k \geq 0$. This also implies that $|\hat{z}_k^i| \leq R$ and concludes the bound on $\bar{\chi}^z$. For obtaining the remaining two bounds, we analyze the x^i, u^i dynamics of agent i , where we introduce the convex conjugate

$$\bar{f}^i(u, z) = \sup_{x \in \mathbb{R}^n} u^T x - f^i(x) - \frac{\rho}{2} |x - z|^2.$$

We note that the supremum is attained for x_{k+1}^i , if $u = -\rho u_k^i$ and $z = \hat{z}_k^i$ in the previous equation. In addition, due to the properties of the convex conjugate, we conclude that $\bar{f}^i(\cdot, z)$ is $1/\rho$ -smooth and $1/(\rho + L)$ -strongly convex. The conjugate subgradient theorem implies,

$$\nabla_u \bar{f}^i(-\rho u_k^i, \hat{z}_k^i) = x_{k+1}^i,$$

which means that the updates for u_k^i can now be expressed as:

$$u_{k+1}^i = u_k^i + \nabla_u \bar{f}^i(-\rho u_k^i, \hat{z}_k^i) - \hat{z}_k^i.$$

By applying Taylor's theorem, we obtain

$$u_{k+1}^i = u_k^i + \nabla_u^2 \bar{f}^i(\nu_k, \hat{z}_k^i)(-\rho u_k^i) - \hat{z}_k^i + \nabla_u \bar{f}^i(0, \hat{z}_k^i),$$

for some $\nu_k \in \mathbb{R}^n$. By leveraging the fact that, as a result of strong convexity and smoothness of $\bar{f}^i(\cdot, z)$, the Hessian $\nabla_u^2 \bar{f}^i(\cdot, z)$ is upper and lower bounded by $1/\rho$ and $1/(\rho + L)$, respectively, we conclude

$$|u_{k+1}^i| \leq \left(1 - \frac{\rho}{\rho + L}\right) |u_k^i| + |\hat{z}_k^i - \nabla_u \bar{f}^i(0, \hat{z}_k^i)|. \quad (54)$$

By a similar argument based on Taylor's theorem, we can bound the last term in the previous equation by

$$|\hat{z}_k^i - \nabla_u \bar{f}^i(0, \hat{z}_k^i)| \leq |x_*^i| + \frac{L}{\rho + L} |\hat{z}_k^i|.$$

Unrolling the recursion in (54) and exploiting the fact that $u_0^i = 0$ yields

$$|u_k^i| \leq \frac{\rho + L}{\rho} |x_*^i| + \sup_{k \geq 0} |\hat{z}_k^i|,$$

where the last term is bounded by R . Finally, due to the $1/\rho$ -smoothness of $\bar{f}^i(\cdot, z)$, we conclude $|x_{k+1}^i| \leq |u_k^i|$, which yields the desired result as follows,

$$\begin{aligned} |\chi^{di}| &= |\alpha x_{k+1}^i + u_k^i| \leq \alpha |x_{k+1}^i| + |u_k^i| \leq (\alpha + 1) |u_k^i| \leq (\alpha + 1) \left(\frac{\rho + L}{\rho} |x_*^i| + \sup_{k \geq 0} |\hat{z}_k^i| \right) \\ &\leq (\alpha + 1) \left(\frac{\rho + L}{\rho} |x_*^i| + R \right). \end{aligned}$$

□

F. Diminishing Communication Threshold

In the main text, we focused our presentation on fixed communication thresholds. However, it is important to note that our approach and our analysis can be easily extended to the case where communication thresholds Δ are varied as a function of the number of iterations. For example, it is straightforward to show that for any vanishing sequence Δ_k , our iterates indeed converge to the minimizer of (3).

Corollary F.1. *Let the assumptions of Thm. D.5 be satisfied and let $\Delta_k \geq 0$ be such that $\Delta_k \rightarrow 0$ for $k \rightarrow \infty$. Then $\lim_{k \rightarrow \infty} |\xi_k - \xi_*|^2 \rightarrow 0$.*

Proof. According to Thm. D.5, the following holds,

$$V_{k+1} \leq \tau^2 V_k + \bar{\sigma}(Q) \Delta_k^2,$$

see (52). We now apply Lemma 3 in Sec. 2.2 in (Polyak, 1987), which yields the desired result. \square

We can also derive an explicit convergence rate. In fact, the following corollary proves that if Δ_k^2 is the form $q/(k+1)^t$, where $q > 0$ and $t > 0$ are constants and k is the iteration number, $|\xi_k - \xi_*|^2$ converges at a rate of $\mathcal{O}(1/k^t)$.

Corollary F.2. *Let the assumptions of Thm. D.5 be satisfied and let $\Delta_k^2 \leq \frac{q}{(k+1)^t}$, $\forall k \geq 0$, $t > 0$. Then, the following holds for all $k \geq 0$:*

$$|\xi_k - \xi_*|^2 \leq \frac{1}{\bar{\sigma}(P)} \left(\frac{k_0}{k + k_0} \right)^t c_0,$$

where $k_0 = \frac{1}{\left(\frac{2}{1+\tau^2}\right)^t - 1}$ and $c_0 = \max \left\{ \frac{2\bar{\sigma}(Q)q}{1-\tau^2}, \bar{\sigma}(P) |\xi_0 - \xi_*|^2 \right\}$.

Proof. According to (52), the following holds,

$$V_{k+1} \leq \tau^2 V_k + \bar{\sigma}(Q) \frac{q}{(k+1)^t}.$$

We make the following claim:

$$V_k \leq c_0 \left(\frac{k_0}{k + k_0} \right)^t, \quad \forall k \geq 0.$$

We prove the claim by induction. The claim holds for $k = 0$ due to the fact that $c_0 \geq \bar{\sigma}(P) |\xi_0 - \xi_*|^2$. We therefore assume that the claim holds for k and show that this implies that the claim holds for $k + 1$. This yields

$$\begin{aligned} V_{k+1} &\leq \tau^2 V_k + \bar{\sigma}(Q) \frac{q}{(k+1)^t} \\ &\leq \tau^2 c_0 \left(\frac{k_0}{k + k_0} \right)^t + \bar{\sigma}(Q) \frac{q}{(k+1)^t} \\ &\leq c_0 \left(\frac{k_0}{k + k_0 + 1} \right)^t \left(\tau^2 \left(\frac{k + k_0 + 1}{k + k_0} \right)^t + \frac{\bar{\sigma}(Q)q}{c_0 k_0^t} \left(\frac{k + k_0 + 1}{k + 1} \right)^t \right) \\ &\leq c_0 \left(\frac{k_0}{k + k_0 + 1} \right)^t \left(\tau^2 \left(\frac{k_0 + 1}{k_0} \right)^t + \frac{\bar{\sigma}(Q)q}{c_0} \left(\frac{k_0 + 1}{k_0} \right)^t \right) \\ &\leq c_0 \left(\frac{k_0}{k + k_0 + 1} \right)^t, \end{aligned}$$

and completes the induction argument. \square

G. Additional Experiments and Hyperparameters

We ran various experiments in order to assess the performance of the event-based distributed learning algorithm (Alg. 1). In the comparative studies, we choose FedAvg (McMahan et al., 2017), FedProx (Li et al., 2020a), SCAFFOLD (Karimireddy et al., 2020) and FedADMM (Zhou & Li, 2023) as baselines, since these methods have been developed to address challenges such as data heterogeneity and communication efficiency. For a fair comparison in terms of computation resources in all setups, each of the agents are run for the same number of local gradient steps.

We include the first example in Sec. 5, which showcases how two image classifiers (for MNIST and CIFAR-10 datasets) can be trained in a distributed and communication efficient way. Our setup included $N = 10$ agents for MNIST, each storing data for a single digit, resulting in the most extreme non-i.i.d. distribution of data among agents. For a CIFAR-10 classifier, the data are distributed among $N = 100$ agents according to a Dirichlet distribution, i.e., we sample $p_a \sim \text{Dir}N(\beta)$, where N is the number of agents and $\beta = 0.5$. We then assign a $p_{a,j}$ proportion of the training data of class a to agent j .

We applied our implementation to train a fully connected neural network on the MNIST dataset and a convolutional network on the CIFAR-10 dataset. The classifier model has 4 convolutional layers, each with 3×3 kernels and 32, 64, 128, and 256 filters, respectively, followed by three fully connected layers with ReLU activation functions. After each set of convolutions, a 2×2 max pooling layer is applied, followed by a ReLU activation. We train the MNIST classifier model using Alg. 1, where we replace the full minimization step of each local objective with five steps of stochastic gradient descent with a learning rate of $l_r = 10^{-1}$, and the CIFAR-10 model with 3 epochs of stochastic gradient descent (batch size 20, learning rate $l_r = 10^{-3}$). Further hyperparameters are listed in Tabs. 3 and 4.

Tab. 1 in Sec. 5 summarizes the main result of this paper, by comparing the performance of different methods. From this table, it is clear that Alg. 1 achieves the same test accuracy with less communication cost. The communication configurations for Tab. 1 are summarized in Tab. 2.

Fig. 8 illustrates the trade-off between accuracy and communication load. The results demonstrate that our event-based approach consistently achieves higher accuracy with fewer communication events compared to baselines. Each point in Fig. 8 represents a different value of Δ , where Δ monotonically increases along the curve, demonstrating that with our algorithm and a well-chosen Δ threshold, communication among agents can be reduced while still achieving a high classification accuracy. Our experimental results indicate that the approach can reduce communication costs by over 30% without significant accuracy degradation. Notably, SCAFFOLD doubles the communication cost due to its dual-package communication protocol. These findings directly translate to the results in Tab. 1 showing total communication events for target accuracies. The extensive experimentation across both small-scale (MNIST) and large-scale (CIFAR-10) scenarios demonstrates the scalability and effectiveness of our event-based communication strategy, particularly for large-scale distributed learning problems.

The next sections provide additional numerical experiments. We first show an example based on LASSO where the local objectives are strongly convex (Sec. G.1 and G.2). In this setup, our theoretical results apply. Sec. G.3 shows how our algorithm can train an MNIST classifier, when only local communications are allowed, as specified by a given communication graph. In such a setup, the baselines FedAvg, FedProx, SCAFFOLD and FedADMM are not applicable.

Algorithm	MNIST Target Accuracy			CIFAR-10 Target Accuracy			
	80%	85%	90%	70%	75%	77%	78%
Alg. 1-randomized ($p_{\text{trig}}, \Delta^d$)	(0.1, 5)	(0.1, 4)	(0.1, 1)	(0.2, 4.5)	(0.1, 3.75)	(0.2, 3.5)	(0.7, 3.75)
Alg. 1-Vanilla (Δ^d)	(3)	(2)	(1)	(4.25)	(3.25)	(3.25)	(1.75)
FedADMM ($part_rate$)	0.4	0.6	1.0	0.4	0.5	0.7	0.9
FedAvg ($part_rate$)	0.4	1.0	-	0.1	-	-	-
FedProx ($part_rate$)	0.5	1.0	-	0.2	-	-	-
SCAFFOLD ($part_rate \times 2$)	0.4×2	0.5×2	0.8×2	0.2×2	-	-	-

Table 2: Communication configurations across algorithms. Values represent the probability of communication for baseline methods. SCAFFOLD values are doubled due to double package transmission per round.

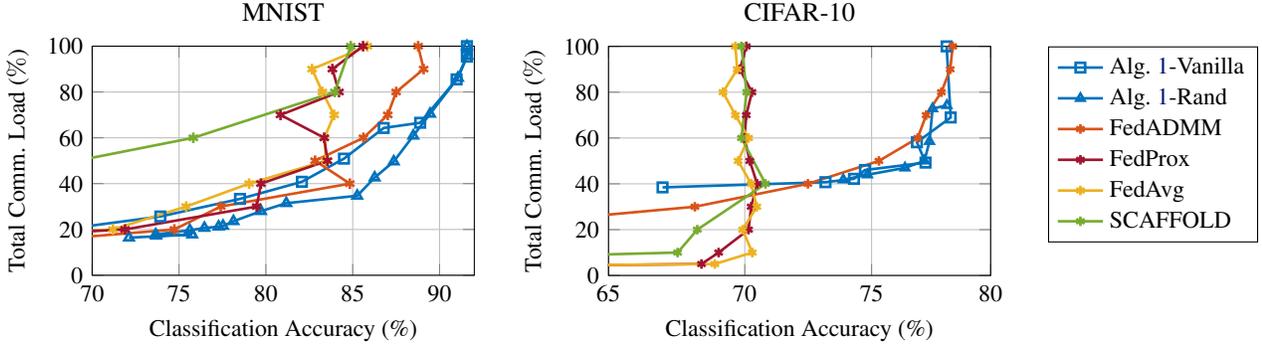


Figure 8: The figure compares different federated learning methods on the MNIST and CIFAR-10 datasets with respect to the resulting trade-off between total communication load and classification accuracy on the test set. In the MNIST case (left), randomization includes agent-to-server communication with 0.1 probability. For CIFAR-10 (right), randomization incorporates server-to-agent communication with 0.2 probability. Points along each curve represent different Δ thresholds, demonstrating the relationship between communication reduction and model accuracy.

Table 3: The table summarizes the hyperparameters used for distributed training of MNIST classifier (Fig. 8, Tab. 1)

Hyperparameter	Value
number of agents (N)	10
size of neural network layers	[400, 200, 10]
learning rate (gradient descent step-size)	0.1
number of iterations	100
$\Delta^d = \Delta, \Delta^z = 0.1 \times \Delta$	range between [0, 10]
μ (FedProx)	0.1
augmented lagrangian parameter (ρ) (FedADMM, Alg. 1)	1
n_g (SCAFFOLD)	1

Table 4: The table summarizes the hyperparameters used for the distributed training of CIFAR-10 classifier (Fig. 8, Tab. 1)

Hyperparameter	Value
number of agents (N)	100
augmented lagrangian parameter (ρ) (FedADMM, Alg. 1)	0.01
learning rate	0.01
momentum	0.9
number of iterations	150
number of local epochs	3
batch size	20
$\Delta^d = \Delta, \Delta^z = 0.01 \times \Delta$	range between [0, 4]
μ (FedProx)	0.1
n_g (SCAFFOLD)	1

G.1. Linear Regression and LASSO with Non-i.i.d. Data

We conduct numerical experiments based on the following distributed learning problem:

$$\begin{aligned}
 \min_{x \in \mathbb{R}^n, z \in \mathbb{R}^n} \quad & \sum_{i=1}^N \frac{1}{2} |A^i x^i - b^i|^2 + \lambda |z|_1, \\
 \text{subject to} \quad & x^i - z = 0, \quad i = 1, \dots, N,
 \end{aligned} \tag{55}$$

where $A^i \in \mathbb{R}^{m \times n}$, $b^i \in \mathbb{R}^m$. In the data generation process, we generate samples from three different distributions: a standard normal distribution, a Student's t distribution with one degree of freedom, and a uniform distribution in the range $[-5, 5]$. These samples are concatenated to form a single dataset, which is then partitioned into subsets for each agent i to

obtain (A^i, b^i) . Finally, we normalize the feature vectors and target values for each agent to prepare the data for the learning problem. In this non-i.i.d. setting, local optimal points of individual agents x_*^i are far away from each other, and their average $\sum_{i=1}^N x_*^i/N$ is also far away from the global optima x_* . The experiments were run for $T_{\max} = 50$ steps, which are required for Alg. 1 to converge to the global optimal point with high accuracy. Fig. 9 illustrates the communication load against the absolute difference between the objective function value f and the optimal value f^* , where the communication load is defined as the number of communications accumulated over time.

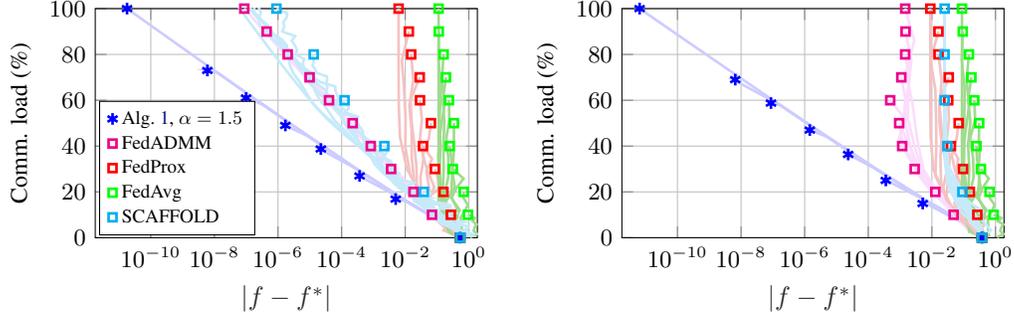


Figure 9: The figure shows the communication load versus accuracy trade-off for the different methods applied to two distinct problems derived from (55): linear regression ($\lambda = 0$, left panel), and on the right, LASSO ($\lambda = 0.1$, right panel).

In the first scenario, we set $\lambda = 0$ to obtain a linear regression problem, where the proposed algorithm with relaxation parameter $\alpha = 1.5$ clearly outperforms baseline methods by a large margin. We note that the gap between the global and local optimal points prevents FedAvg and FedProx from converging to the optimal point f^* .

For the second case, we set $\lambda = 0.1$ to solve the LASSO problem. By assumption, FedAvg, SCAFFOLD and FedADMM require the local objective functions to be smooth. However, we allow handling nonsmooth local objective functions, which is relevant to the distributed learning problems with ℓ^1 regularization. To avoid a noncontinuous gradient for the local minimization for SCAFFOLD, FedADMM, FedAvg and FedProx, the local update step is carried out by the following local gradient,

$$\nabla_{x^i} \tilde{f}^i(x^i) = A^{i\top} (A^i x^i - b^i) + \frac{\lambda}{N} \begin{cases} \text{sgn}(x^i) & |x^i| > \delta \\ \frac{1}{\delta} x^i & |x^i| \leq \delta \end{cases}, \quad (56)$$

where δ can be chosen as small as $1e - 16$ (double precision machine epsilon). However, we found that the results are largely unaffected by the choice of δ .

Table 5: The table summarizes the hyperparameters used for the distributed linear regression and LASSO experiments (Fig. 9).

Hyperparameter	Value
number of agents (N)	50
augmented lagrangian parameter (ρ)	1
gradient descent step-size	1
number of iterations	50
$\Delta^d = \Delta^z = \Delta$	range between $[0, 10^{-2}]$

G.2. Effect of Communication Drops

To observe the effect of communication drops, we repeated the same LASSO experiment in (55) with hyperparameters in Tab. 6, but this time, we allow the transmission of information from the agents to the server to fail with a probability of 0.3. As seen in the second panel of Fig. 10, if we have no reset, i.e., $T = \infty$, the algorithm cannot converge and a significant error remains. On the left panel, the trade-off between communication load and suboptimality is presented. More frequent reset operations lead to a faster convergence and less error, in exchange for additional communication cost that comes with the reset.

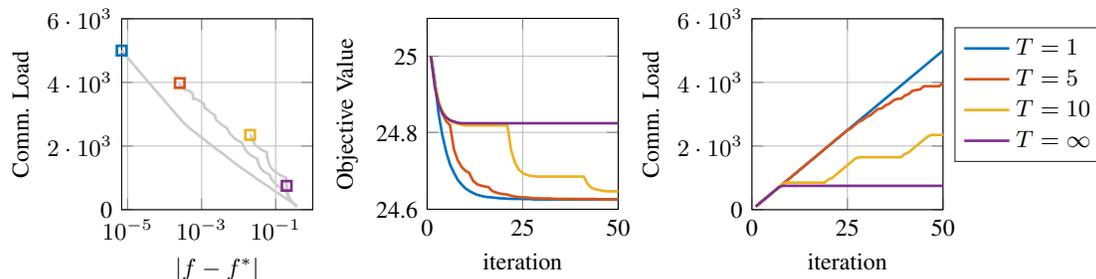


Figure 10: The left panel presents the trajectory of communication load versus suboptimality of the objective function value. The panel in the center shows the evolution of the objective function for different values of the reset period for a drop rate of 0.3 and for $\Delta = 10^{-3}$, whereas the right panel shows the cumulative communication load over time in addition to the reset communication at each T step.

Table 6: The table summarizes the hyperparameters used for the distributed LASSO experiment against communication drops (Fig. 10).

Hyperparameter	Value
number of agents (N)	50
L1 regularization parameter (λ)	0.1
augmented lagrangian parameter (ρ)	1
relaxation parameter (α)	1
gradient descent step-size	1
number of iterations	50
$\Delta^d = \Delta^z = \Delta$	10^{-3}

G.3. Distributed Training on a Graph

Our distributed learning algorithm, Alg. 2, is general enough to train a machine learning classifier over a network of agents; the network structure can be encoded by a proper selection of the linear constraint matrices A and B (see App. A for further details). Our framework therefore generalizes well beyond server-client structures, and our theoretical analysis also captures the influence of the network structure on the resulting convergence rate.

In order to highlight the versatility, we train an MNIST Classifier over a network of agents. We use a multi-layer perceptron that has the same structure as in Sec. 5 and consider a situation where each agent has only access to the training data of a single digit. Fig. 11 shows the resulting communication load and classification accuracy trade-off on the entire dataset (left), whereas the diagram on the right shows the network structure (only communication along the edges of the graph is allowed). The error bars indicate the range (minimum and maximum) of the classification accuracy among the different agents.

The results shown in Fig. 11 and highlight that a purely random selection of agents (suggested in (Yu & Freris, 2023)) results in a worse trade-off curve, which further motivates our event-based strategy. We also apply our algorithm to a much larger distributed learning problem with 50 agents and where the corresponding accuracy versus communication trade-off is shown in Fig. 12, together with the agent network that has been used.

Table 8: The table summarizes the hyperparameters used for the distributed linear regression experiment over a graph (Fig. 12).

Hyperparameter	Value
number of agents (N)	50
augmented lagrangian parameter (ρ)	10^{-5}
number of iterations	17×10^3
Δ^x	range between $[0, 1]$

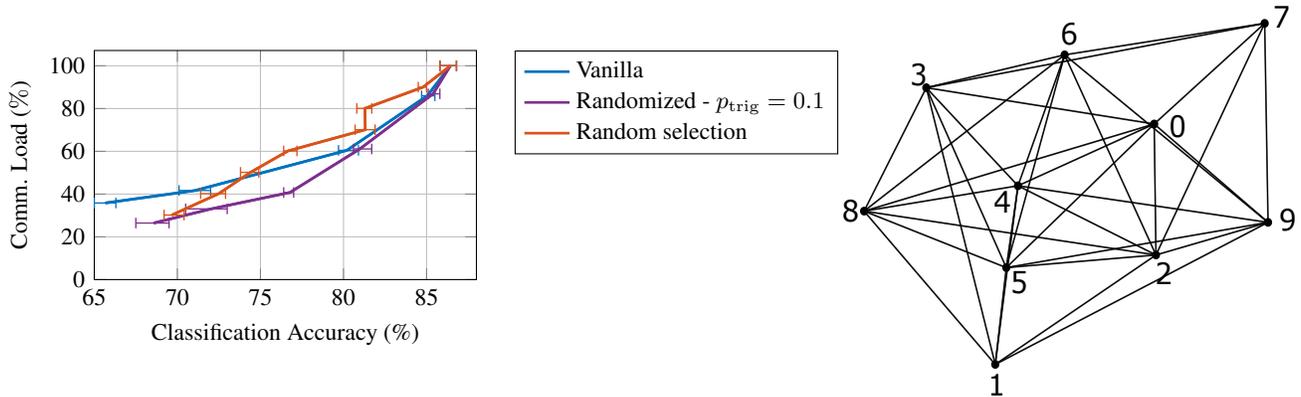


Figure 11: The figure shows a comparison of the vanilla event-based and the randomized event-based communication strategy (see Sec. 2) with a purely random selection of agents. The outcome of a purely random strategy is consistently worse with respect to the resulting trade-off between communication load and classification accuracy. The right panel visualizes the agent network with ten agents connected with 70 edges.

Table 7: The table summarizes the hyperparameters used for the distributed training of MNIST classifier over a graph (Fig. 11).

Hyperparameter	Value
number of agents (N)	10
size of neural network layers	[400, 200, 10]
learning rate (gradient descent step-size)	5×10^{-3}
augmented lagrangian parameter (ρ)	5×10^{-3}
number of iterations	10^3
number of gradient steps per iteration	5
Δ^x	range between [0.0, 0.2]

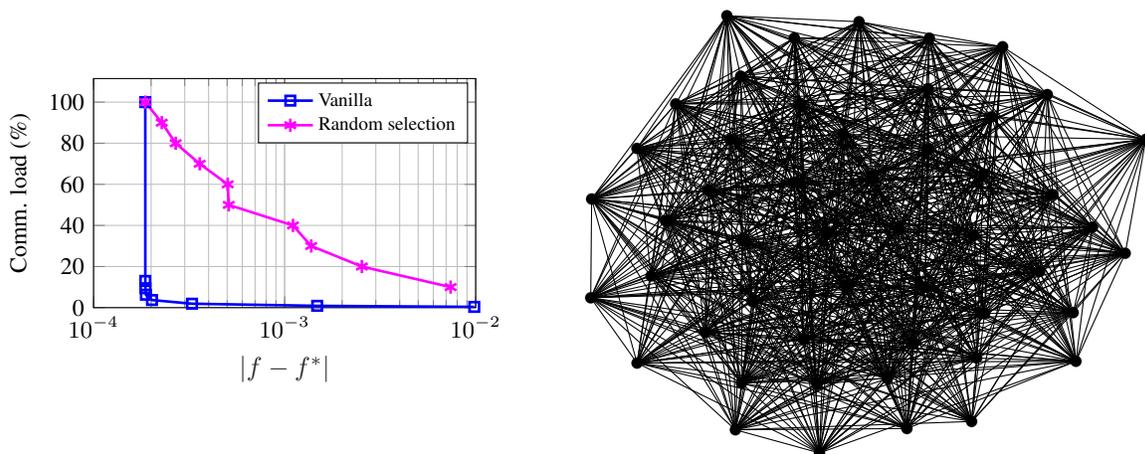


Figure 12: The first panel shows the comparison of the communication load versus solution accuracy for different communication methods applied to the linear regression problem derived from (55) ($\lambda = 0$). The right panel visualizes the agent network with 50 agents connected with 1762 edges.