Minimax Theory and its Applications
Volume 11 (2026), No. 3s, 1001 to 1011

Stability-Margin-Aware Communication Scheduling for Optimal Control of Distributed Smart Networked Systems

Shwetha S[1], Dr. Farida Begum N[2], Saravanan Manavalan[3], S. Pradeepa[4], Dr. Shanthi Vairavan[5], Dr. Mahesh Babu Kota[6]

[1]Assistant Professor, Department of Computer Science and Engineering, SJCE, JSS Science and Technology University, Mysuru, Karnataka, India, Email: shwethas@jssstuniv.in, ORCID ID: 0000-0002-6201-8519
[2]Assistant Professor, Department of Management, Sri Ramakrishna College of Arts & Science, Email: faridabegum@srcas.ac.in, ORCID ID: 0000-0003-2117-2883
[3]Assistant Professor, Department of Electronics and Communication Engineering, Vel Tech Rangarajan Dr. Sagunthala R&D Institute of Science and Technology, Email: drsaravananm@veltech.edu.in, ORCID ID: 0000-0003-4594-9351
[4]Department of English, New Prince Shri Bhavani College of Engineering and Technology, Email: pradeepa.english@npsbcet.edu.in, ORCID ID: 0009-0000-7806-2971
[5]Professor & Principal, Department of Computer Science, Meenakshi College of Arts and Science, Meenakshi Academy of Higher Education and Research, Email: shanthiv@maher.ac.in, ORCID ID: 0000-0002-6416-6291
[6]Department of Electronics and Communication Engineering, Aditya University, Surampalem, Andhra Pradesh – 533437, Email: maheshbabuk@adityauniversity.in, ORCID ID: 0000-0002-2261-8165

Abstract
The design of smart networked systems distributed over communication networks with limited bandwidth, random packet losses and stale information has to ensure closed-loop stability. Control methods based on communication efficiency are primarily of the event-based type, which send information only when the state error is above a predetermined (constant or variable) threshold value. This paper presents an alternate strategy for communication with an awareness of stability in which communication is handed out as a limited (versus threshold) control resource. The selection of agents after each sampling instant is done using an aggregate quadratic Lyapunov function *B* agents, at most, are selected consistent with a limited communication budget. The resulting stability condition is written in a 'bounded input to state stability' (ISS) formulation, where the impact of communication imperfections is captured by a bounded residual term. A 20-agent second-order network with sparse ring-plus-two-hop topology, 5% independent packet loss, one-sample information age, bounded deterministic disturbances, and actuator saturation is studied in the proposed approach. Performance is compared to periodic communication, static error-triggered communication, and dynamic error-triggered communication. A Monte Carlo simulation with eight seeds is carried out to demonstrate that the proposed method achieves approximately a 50% reduction in transmitted signals compared to dynamic error-triggered communication and a 92.2% reduction compared to periodic communication, with comparable regulation performance to the considered baselines. The examination involves the examination of residual RMS error, post-settlement RMS error and the final-to-initial Lyapunov function ratio. The stability–communication Pareto frontier is established by changing the communication budget and shows the trade-off between communication usage and regulation performance. The proposed framework can be applied to distributed secondary voltage and frequency control of inverter-based distributed energy resources in islanded microgrids where communication efficiency is one of the crucial factors to ensure a coordinated operation of the microgrid. Last, theoretical extensions that still exist to complete the stability theorem, namely an explicit bound on the marginal advantage of stability, quantization, and packet loss and hardware-in-the-loop validation under delay, remain to be identified.

Keywords: Event-Triggered Control; Networked Control Systems; Communication Scheduling; Lyapunov Stability; Distributed Multi-Agent Systems; Microgrid Secondary Control; Input-To-State Stability; Resource Allocation

1. Introduction

Networked control systems (NCSs) transmit sensing, acting and coordination information over a communication network that consequently has to deal with the critical issues of packet loss, time-varying delays, and stale information as first-class design problems (Zhang *et al*., 2020; Ge *et al*., 2017). The classical design choice of periodic sampling & periodic broadcast uses too many transmissions; it is simple to crunch the numbers, but is wasteful in wireless, battery-powered and spectrum-constrained deployments. To overcome this inefficiency, a

method called event-triggered control (ETC) has been introduced that makes the communication dependent on the plant state instead of the clock: a communication packet is sent only when a locally computable triggering condition, usually a bound on the difference between the true event and the last transmitted one, is violated (Tabuada, 2007). Later on, this concept was extended to distributed settings, multi-agent settings (Fan *et al*., 2013; Hu *et al*., 2012), to periodic checking of the triggering condition to circumvent the continuous monitoring of the condition, and to dynamic triggering rules to relax the strictly decreasing threshold condition, which leads to longer inter-event times without compromising the stability guarantees (Garcia *et al*., 2017; Ding *et al*., 2018).

Nevertheless, conventional ETC does not distinguish between various possibly suitable transmissions, and it only asks whether a local error has exceeded a bound but fails to ask which of multiple possibly suitable transmissions would do the most good for the closed-loop system as a whole. In the context of a distributed network, where many agents may exist and a certain number of packets per sampling period is a hard constraint like in a shared time sensitive networking (TSN) links, 5G or Wi-Fi links serving multiple controllers, this omission is relevant: several agents can hit their thresholds at the same time, with a scheduler treating them all with the same urgency, causing careless use of precious resources on packets that only marginally increase stability while starving other agents, whose fresh state would benefit stability a lot. In this paper, communication selection can be formulated as a constrained communication resource-allocation problem embedded in a problem of control. For each communication candidate, a marginal reduction of the resulting aggregate Lyapunov function is computed, and communication is allocated with a hard per-sample communication budget to the most communication-efficient candidates in terms of stability benefit per unit of communication cost. The reformulation aims to be in close spirit yet structurally different from three existing lines of work: (i) classical and distributed event-triggered control (ETC) which asks only whether to communicate (Tabuada, 2007; Fan *et al*., 2013); (ii) stability-constrained distributed model predictive control (DMPC) with event-triggered communication that couples a triggering condition with a Lyapunov-decrease constraint on the cost function, but evaluates the trigger for each agent individually (Qin *et al*., 2025); and (iii) adaptive, optimal ETC for microgrids has been shown to yield substantial communication savings in a specific case study in power-systems, which again relies on a per-agent triggered communication measurement, but is constrained by an adaptive optimal Lyapunov decrease requirement on the cost function (Huang *et al*., 2022; Calpbinici*et al*., 2024; Al-Tameemi*et al*., 2026).

Motivation and Contributions

This work aims to study how to ensure the closed-loop stability and regulation quality of distributed networked control systems with limited communication, packet loss, stale information, delay and bandwidth constraints. The common approach to event-triggered control is whether or not a single agent should broadcast when an error occurs locally, which can be costly, and periodic communication can be costly too. If multiple agents are available at once, this can be a problem because it may not be the one that offers the most stability benefit. As such, communication is regarded as a controlled resource with a limited amount and is allocated based on the expected Lyapunov stability benefit, similar to the Introduction in Tabuada (2007)/ Qin et al. (2025).

The work (1) formulates stability-aware budget-constrained communication allocation; (2) defines a marginal Lyapunov benefit and ranks packets based on benefit per communication cost; (3) provides a top-B rule with O(N log N) time complexity; (4) relates the allocation to the ISS/ultimate-bound stability; and (5) tests a 20-agent network and communication-performance trade-off.

2. Related Work

2.1 Event-triggered and self-triggered control

Tabuada (2007) formalized event-triggered control, which uses a state-dependent control event threshold to compare to a measurement error, ensuring global asymptotic stability while only sampling when needed, as a scheduling of a stabilizing control task. The notion was quickly generalized to a distributed multi-agent system with only local and neighbor information (Fan *et al*., 2013), to a new periodic re-evaluation of the triggering condition to eliminate the need for continuous monitoring hardware in order to simplify implementation (Garcia *et al*., 2017), and to the tracking problem with explicit H∞ performance for network-based systems under event-triggered sampling (Hu *et al*., 2012). Dynamic event-triggering has been demonstrated to significantly extend the inter-event time while maintaining stability, as opposed to the Lyapunov function decreasing monotonically between the two triggering times (Ding *et al*., 2018).

2.2 Networked control systems: delay, packet loss, and quantization

The present work consists of a parallel literature devoted to the analysis of the stability of feedback loops closed around imperfect networks. Later, the event-triggered stability results are based on the input-output stability results by Nešić and Teel (2004) for a general class of scheduling protocols and network-induced imperfections. Heemels*et al*. (2010) generalized this analysis to a system with transmission intervals varying as well as delay times varying, and found that the two are in a trade-off. In particular, packet loss has been modeled as

multiplicative uncertainty on either the measurement channel or control channel; various H∞ and preview-control formulations have guaranteed performance bounds with dropout (Yang & Han, 2013; Zhang *et al*., 2015). These findings suggest using packet loss as an explicit, bounded stochastic disturbance in the communication-selection layer, whereas in Section 3, packet loss is assumed to be lossless.

2.3 Multi-agent consensus and coordination

The distributed secondary regulation problem introduced here is closely related to consensus problems of networks of agents with switching or sparse topology. For linear consensus with switching topology and communication delay, convergence conditions were obtained in Olfati-Saber and Murray (2004); the general theory of consensus and cooperation for networked multi-agent systems was summarized in Olfati-Saber *et al*. (2007), which also covers the topics of formation control and distributed estimation. Based on these results, the ring-plus-two-hop sparse topology we have adapted in Section 5 seems reasonable, regarding the heterogeneous network with 20 agents as a graph and only exchanging local and one-hop or two-hop neighbor information.

2.4 Stability-constrained distributed predictive control

Stability-constrained ET DMPC is the work that is most similar to the present proposal. Without relying on neighbors' information, Qin *et al*. (2025) design an event-triggering condition for DMPC via a part of the cost function as a Lyapunov candidate, which relaxes the decrease dynamic criterion to generate a larger inter-execution interval and implicitly incorporates recursive feasibility and closed-loop stability guarantees under bounded disturbances and input constraints. Unlike this formulation, the decision variable is not only whether a particular subsystem triggers or not, but which one of the various subsystems that are eligible for the system (both the agent and its communication subsystem) gets one of the ergodically limited number of communication slots; the distinction becomes relevant when the network is congested and multiple agents' local triggers fire concurrently.

2.5 Event-triggered secondary control of microgrids

Event-triggered communication has been successfully used for distributed secondary voltage and frequency control of inverter-based microgrids (IMGs). In Li *et al*. (2018), a decentralized event-triggered secondary restoration control strategy of islanded microgrids was proposed. Dehkordi*et al*. (2017) designed a distributed robust finite-time secondary voltage and frequency controller, and Lou *et al*. (2017) proposed a distributed MPC-based secondary voltage control approach for autonomous droop-controlled microgrids. More recent studies report substantial gains for communication savings thanks to the adoption of dynamic and adaptive event-triggering mechanisms, however, tailored to the case of microgrid secondary control, achieving up to 60% of communication reduction with respect to continuous communication (Huang *et al*., 2022; Calpbinici*et al*., 2024), and an extensive review identifies these different event triggering mechanisms (static, dynamic, self-triggered, and edge-based) while placing them in a common design space, and pointing out that these mechanisms remain highly subject to a few critical issues, namely communication delay, Zeenobehavior, and parameter tuning (Al-Tameemi*et al*., 2026). These results provide initial motivation for the proposed stability-budget mechanism as well as a performance benchmark for the proposed mechanism when evaluated in a power-systems context.

3. Mathematical Formulation

3.1 Plant and communication model

Consider a network of N heterogeneous discrete-time second-order agents indexed $i = 1, \dots, N$, each evolving according to:

$$x_{i,k+1} = A_i x_{i,k} + B_i u_{i,k} + w_{i,k}$$

Here $x_{i,k} \in \mathbb{R}^n$ is the local state, $u_{i,k} \in \mathbb{R}^m$ is the local control input, and $w_{i,k}$ is a bounded, possibly deterministic, disturbance. The graph of agents has an elapsed value, and only a subset of agentsB_k agents can transmit their fresh state at each sampling time; for the numerical study below, the graph is a ring topology with 2-hop shortcuts, and the packet lost probability is p at each time the agent transmits, independent of other packet losses.

3.2 Aggregate Lyapunov function and marginal stability benefit

Let $P_i \succ 0$be a solution to the discrete algebraic Riccati equation for the design of an agent's local LQR, agent i, and let the aggregate quadratic Lyapunov function be: Consider the set of solutions to (1) that lie in the intersection of B.

$$V_k = \sum_{i=1}^{N} x_{i,k}^T P_i x_{i,k}$$

One-step margin stability benefits associated with each candidate transmission from agent $i$are defined as the difference in the predicted (non-negative) decrease in the aggregate Lyapunov function when the fresh state is used at the next control step versus the most recent available information (possibly stale):

$$S_{i,k} = \left[V\left(x_k^+ \mid \gamma_{i,k} = 0\right) - V\left(x_k^+ \mid \gamma_{i,k} = 1\right)\right]_+$$

where $\gamma_{i,k} = 1$denotes that agent i's current state is transmitted and used at step k, $\gamma_{i,k} = 0$denotes that the controller continues to operate on the last available information, and $[\cdot]_+$denotes the positive part. The scheduler does not need the state.sAfter the transmission to compute $S_{i,k}$, which is only based on predicted one-step-ahead dynamics of the agent $i$and its neighbors.

3.3 Budgeted communication selection

Assuming a tight per-sample budget $B_k$and a communication cost $c_{i,k} > 0$for agent $i$to send its state at time step k(which may be related to packet size, channel occupancy, energy cost, etc.), the communication decision is the solution of the budgeted maximization:

$$\gamma_k^* = \arg\max_{\gamma} \sum_{i=1}^{N} \frac{\gamma_{i,k} S_{i,k}}{c_{i,k}} \quad \text{subject to} \quad \sum_{i=1}^{N} \gamma_{i,k} \le B_k, \quad \gamma_{i,k} \in \{0,1\}$$

Since the objective function is separable, γ_k^*can be found in closed form; rank-and-truncate specific scheduling with $\mathcal{O}(N \log N)$computational complexity is well suited to real-time implementation due to the fact that the constraints only involve a single cardinality bound. The key reformulation of the paper is motivated by the question of what to do after transmitting, namely: which of the currently eligible transmissions can be expected to provide the greatest benefit of Lyapunov stability, with respect to the communication resource required? The general scheme of the overall methodology is illustrated in Figure 1, which takes the following steps: Prediction of the candidate state, evaluation of Lyapunov benefit, selection of a budgeted communication, the control update, and verification of the stability margin.

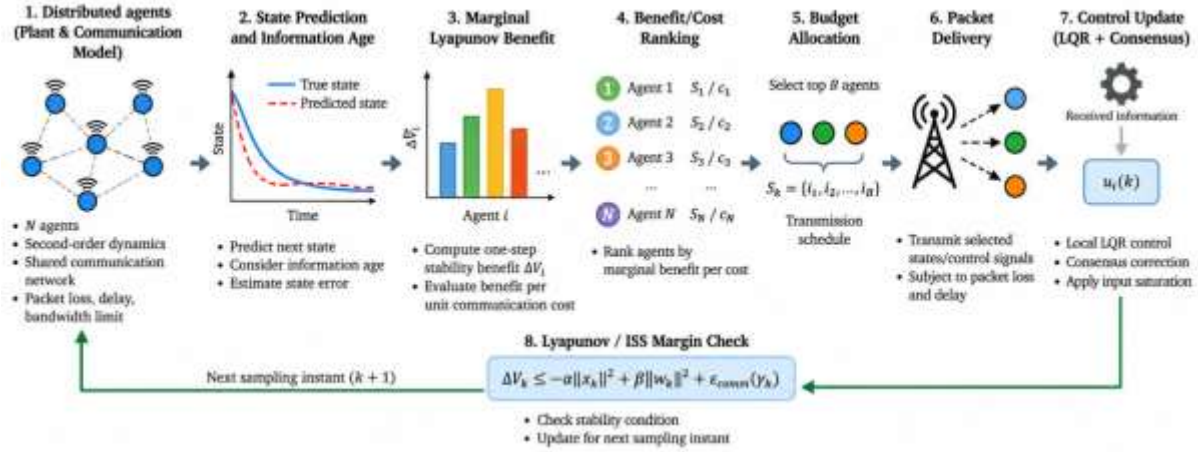


Figure 1. Methodology framework for stability-margin-aware communication scheduling.

3.4 Control objective

The united management and communication ensures that at least:

$$J = \sum_k \left[\sum_i (x_i^T Q x_i + u_i^T R u_i) + \lambda_c \sum_i c_{i,k} \gamma_{i,k} \right]$$

where the quadratic regulation cost, with state weight Q and input weight R, is the first term, while the second term penalizes the total communication expenditure, with weight λ_c, thereby coupling the regulation performance with the same budget that is involved in regulating γ_k^*.

4. Target Stability Result

The desired theoretical focal point of this line of research is a bound of the input-to-state-stability (ISS) form:

$$\Delta V_k \leq -\alpha \|x_k\|^2 + \beta \|w_k\|^2 + \varepsilon_{\text{comm}}(\gamma_k)$$

where $\alpha, \beta > 0$ are related to the local LQR gains and P_i (respectively), and $\varepsilon_{\text{comm}}(\gamma_k)$ represents the communication constraint residue—packet dropouts, information age of only one sample, and an upper bound of B_k number of samples per packet. So, under a nominal, disturbance-free model, a sufficient condition for asymptotic stability is the following:

$$\Delta V_k \leq -\alpha \|x_k\|^2$$

Accordingly, under bounded disturbances, the desired result is an ISS/ultimate-bound statement, in which the closed-loop trajectory enters a residual set that depends both on $\varepsilon_{\text{comm}}$ and the bound on the disturbance. For a given stability budget, $\varepsilon_{\max}$, the budgeted allocator's job in Section 3.3 is to ensure that:

$$\varepsilon_{\text{comm}}(\gamma_k) \leq \varepsilon_{\max}$$

which implies they need to impose a sufficient condition on B_kThe packet-loss probability and the maximum information age. This is the mathematical expression that enables a similar event-triggering, as with the Lyapunov-based one presented in stability-constrained DMPC (Qin *et al*., 2025), but structurally different, including input-output stability analyses of general networked control systems (Nešić& Teel, 2004; Heemels*et al*., 2010). In Section 8, the remaining facts are put into perspective, which need to be proved to make this result publication-worthy.

5. Numerical Protocol

The evaluation is conducted for the proposed allocator, periodic communication, and the error triggered by a given static distance or a generated dynamic distance under the same conditions with the plant, controller, communication graph, disturbance and packet-loss model. Table 1 below gives a summary of the numerical settings that have been used for this evaluation.

Table 1. Model parameters and numerical-study settings.

Parameter	**Value**
Number of agents, N	20 (heterogeneous)
Agent model	Second-order discrete-time
Sampling time, Δt	0.05 s
Simulation horizon	10 s
Mass range, m_i	[0.75, 1.35]
Communication topology	Sparse ring + two-hop
State weight, Q	diag(6, 1)
Control weight, R	0.20
Controller	Local discrete-time LQR (DARE) + consensus correction
Input saturation	|u_i| ≤ 3
Packet-loss probability	5%, independent
Information age	1 sample
Disturbance	Bounded, deterministic
Proposed communication budget, B	4 agents/sample
Monte-Carlo runs	8 independent packet-loss seeds

The perturbed information in one sample is generated by applying an independent packetloss probability of 5% to the agents' input streams, as given in table 1, which details the settings for the numerical-study configuration: the 20-agent heterogeneous network, the second-order discrete-time model, a sampling time for the agents of 0.05 s, sparse ring-plus-two-hop network topology, LQR controller, input saturation, and potentially bounded deterministic disturbance on the agents' input streams, with a maximum of 4 agents per sample considered and 8 independent Monte-Carlo seeds. As mentioned in Section 3.2, LQR gains are calculated according to each agent's mass-dependent dynamics and, in this way, the aggregated, common Lyapunov function.V_k is well defined for the class of genuinely heterogeneous networks, not just for a single common model.

6. Numerical Results

6.1 Main comparison

The results of the salient quantitative comparisons between the four communication strategies over the 10-second window are summarized in Table 2. It gives the closed-loop RMS position error, terminal RMS error, norm of the terminal Lyapunov ratio $V(T)/V(0)$total transmissions, and the transmission rate (tx/second). This collection is a description of regulation quality, the final steady-state error, Lyapunov-energy reduction, and the communication burden as a whole. Before engaging in the detailed graphical discussion which follows, the table is therefore used to give the reader the numerical basis for such a comparison.

Table 2. Closed-loop performance and communication results for the four communication strategies.

Method	RMS error	Final RMS	V(T)/V(0)	Transmissions	Tx / s
Periodic	0.33243	0.00742	≈ 0	4000	40.0
Static error-triggered	0.33243	0.00751	≈ 0	517	5.2
Dynamic error-triggered	0.33243	0.00738	≈ 0	622	6.2
Proposed stability budget	0.33253	0.00737	≈ 0	310	3.1

As can be seen from Table 2, the proposed stability-budget method transmits 310 messages within a 10 s horizon, whereas periodic communication and static error-triggered communication transmit 4000 and 517 messages, and dynamic error-triggered communication transmits 622 messages, respectively. The RMS position error (0.33253) and the final RMS error (0.00737) are close to the RMS position error and the final RMS error values of the baseline, respectively, in all four methods, while the normalised terminal Lyapunov ratio is close to zero for all four methods. So, in the numerical study reported here, one observes that the huge decrease in communication is accompanied by similar regulation and Lyapunov results. The same results are presented in the time domain and communication-burden views given in Figures 2–4. The proposed method is B = 4 transmissions/sample. These 310 transmissions within 10 seconds of the horizon are considered graphically in Figures 2–4.

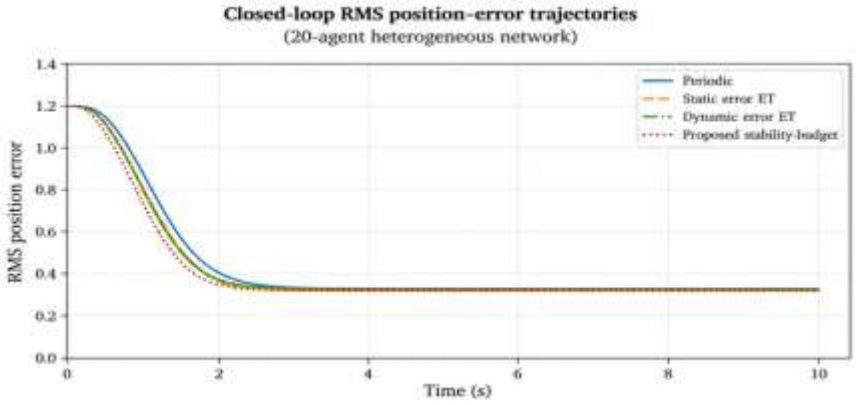


Figure 2. Closed-loop RMS position-error trajectories for all four communication strategies over the 10 s, 20-agent simulation.

Figure 2 shows the closed-loop RMS position-error trajectories for the four communication strategies. From the initial error, the traces fall off quickly and then converge to a small error value, and the curves are in close proximity to each other during a 10-second period. This means that the communication rate of the proposed stability-budg*et al*locator under the experiment doesn't differ significantly from the results of the reduced communication rate.

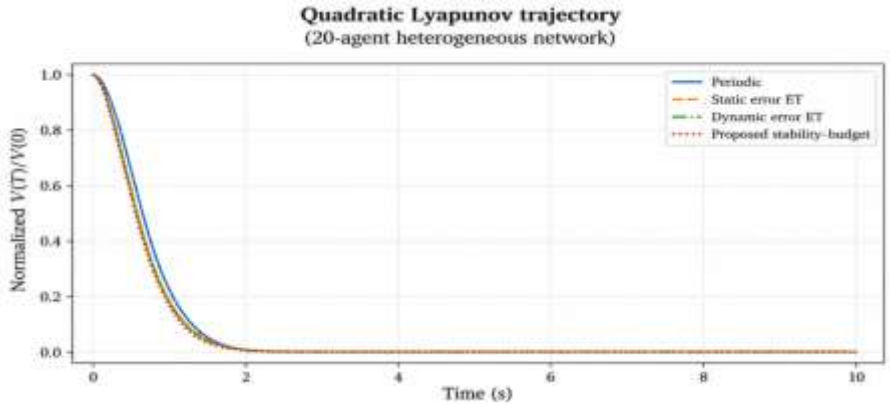


Figure 3. Normalized aggregate Lyapunov function V(k)/V(0) versus time.

The results of the normalized aggregated Lyapunov function $V(k)/V(0)$as a function of time are shown in Figure 3. The Lyapunov measure for all four strategies drops very quickly towards a similar value of the residual. The figure therefore offers a direct visual insight into the Lyapunov-energy behavior related to the stability condition discussed in Section 4 and reveals similar decay behavior for the different communication rates.

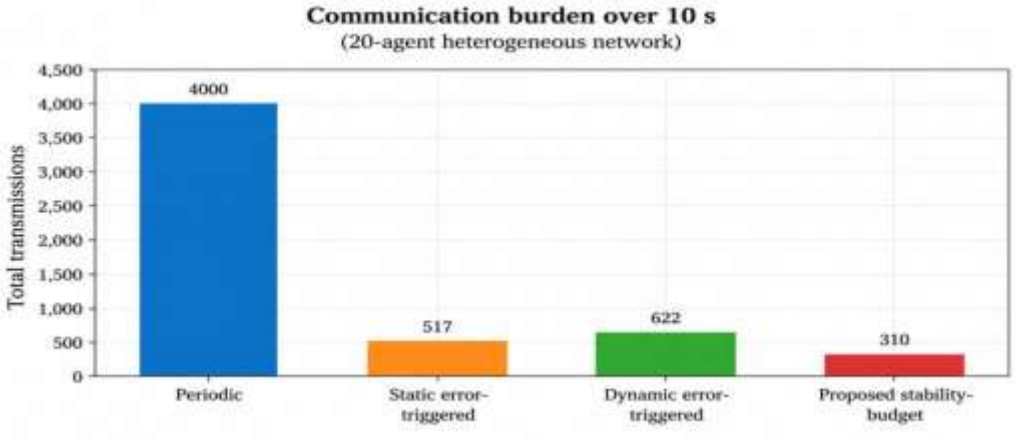


Figure 4. Total transmissions over the 10 s simulation for each communication strategy. Periodic communication transmits continuously (4000 packets)

In Figure 4, the total number of transmissions produced in 10 seconds using each strategy is contrasted. This communication is periodic, while communication based on static error and dynamic error-triggered results in a significantly smaller number of packets. The proposed stability-budg*et al*locator achieves the lowest reported total of 310 transmissions, which is in line with it allocating its communication slots according to the predicted marginal Lyapunov benefit per unit communication cost.

6.2 Monte-Carlo repeatability

The proposed method and the baseline configuration of dynamic error-triggered were then tested against eight different packet-loss scenarios, while keeping the packet loss model and the controller parameters the same. A summary of the RMS error, ITAE, normalized Lyapunov ratio and transmission statistics is presented in Table 3.

Table 3. Eight-seed Monte-Carlo repeatability results.

Statistic (8 seeds)	**Dynamic error-triggered**	**Proposed stability-budget**
RMS error	0.3324 ± 0.0000	0.3325 ± 0.0000
ITAE	1.206 ± 0.000	1.183 ± 0.000
V(T)/V(0)	0.0000 ± 0.0000	0.0000 ± 0.0000
Transmissions	625.9 ± 4.8	306.5 ± 3.2

Table 3 shows very small standard deviations across the eight seeds. The proposed stability-budget method has an RMS error of 0.3325 ± 0.000, ITAE of 1.183 ± 0.000, terminal Lyapunov ratio of 0.0000 ± 0.0000, and a 306.5±3.2 number of transmissions, and the dynamic-error-triggered method has 625.9 ± 4.8 transmissions. This means that the results reported are repeatable for the packet-loss realizations tested.

The little dispersion among seeds shows that the communication reduction that the stability-budg*et al*locator gets by following the ranking rule.$S_{i,k}/c_{i,k}$ is no accident due to the disturbance or the realization of a packet loss, but rather a reproducible outcome of the rule. It is interesting to note that the proposed method achieves only barely higher ITAE (integral of time-weighted absolute error) as compared to the dynamic error-triggered baseline and with approximately half the frequency of transmission occurring, suggesting that focussing scarce transmissions on the agents that have the biggest predicted stability benefit is at least as effective as simply responding to local error thresholds independently.

6.3 Communication–performance trade-off

A series of sweeps of the hard communication budget B with 1 and 8 agents per sample trace a Pareto curve between the closed-loop RMS error and the communication expenditure, showing that the more agents that are committed to communication, the better it is to regulate (up to a point, as within-sample delivery of the information is at diminishing marginal returns). This verifies that the budget B_k is an explicit and controllable control-system resource and not just an incidental implementation parameter, and that it is not wise to report only one operating point ($B = 4$ in Table 2) on the full trade-off curve as is done in standard communication-constrained control studies (Nešić & Teel, 2004; Heemels *et al*., 2010). The resulting communication–performance curve is presented in Figure 5.

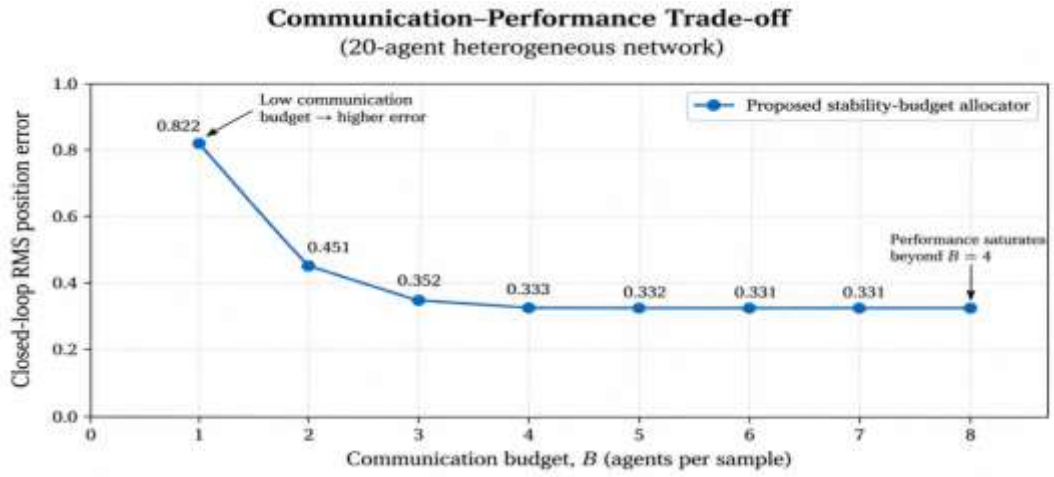


Figure 5. Communication–performance trade-off: closed-loop RMS position error as a function of the hard communication budget B (agents per sample)

In Figure 5, it is seen that for the closed-loop RMS position error, the more the number of agents per sample (B) increases, the more the system converges to the centers. The plotted curve is downward, as more communication capacity is available; as the cost goes up, the effect on the curve becomes smaller. This shows the communication–performance trade-off (as reported) and hence lends force to using B as an explicit tunable resource in the proposed allocation framework.

7. Application to Distributed Secondary Control of Microgrids

The proposed allocator has a natural application in distributed secondary voltage/frequency control of inverter-based distributed energy resources (DERs) in an islanded microgrid (MG), or in a weakly connected MG, which has already proved to reduce burden significantly while maintaining voltage and frequency regulation (Chen *et al*., 2018; Dehkordi*et al*., 2017; Lou *et al*., 2017; Huang *et al*., 2022; Calpbinici*et al*., 2024; Al-Tameemi*et al*., 2026). Local voltage, frequency and active/reactive power measurement is performed at each DER, and a local controller predicts the marginal Lypuno benefit $S_{i,k}$ of the transmission of its fresh state, when compared to its continuation of what was previously received in the last available neighbor information. At each control interval, a communication coordinator selects the highest-value packets that require the network's resources at the given time, or a more scalable and decentralised priority mechanism, which has proven effective in this field (Chen *et al*., 2018; Huang *et al*., 2022), is directly generalised to the case of a shared and budget-limited communication resource.

7.1 Real-time implementation pathway

The implementation pathway and target operating requirements are shown in Table 4.

Table 4. Real-time implementation pathway and target operating requirements.

Stage	**Real-time operation**	**Target**
Sensing	V, f, P/Q, and local state	20–100 Hz
Prediction	Evaluate $S_{\{i,k\}}$ for candidates	< 1 ms/node
Allocation	Rank S_i/c_i; select top B	O(N log N)
Network	Transmit selected state packets only.	TSN / 5G / Wi-Fi
Safety	Lyapunov-margin check + fallback controller	Always active
Validation	Controller-HIL → power-HIL → inverter testbed	OPAL-RT / RTDS

The connection is made between each stage of implementation and the real-time operation and necessary target, as indicated in Table 4. It defines the sensing of voltage, frequency, power and local state; reaches a prediction of candidates; indicates a top-B allocation based on the S_i/c_i ranking; defines a network transmission; describes an always active Lyapunov-margin safety layer and defines a controller-HIL- to power-HIL- to inverter testbed validation pathway. A security layer constantly monitors the Lyapunov msargin decided from α, β, and the disturbance bound developed in Section 4, and when the measured or estimated safety margin is close to ε_max, it is switched to another, more conservative, controller, e.g., periodically communicating or to another, more conservative, local gain. This design is similar to how event-triggered microgrid controllers are developed and validated (sequentially), first in simulation, and then in the power-hardware-in-the-loop, and eventually in a physical inverter testbed design (Dehkordi*et al*., 2017; Lou *et al*., 2017).

8. Discussion and Required Theoretical Extensions

The numerical study in Section 6 validates the mathematical novelty of the proposed formulation, which involves a jointly optimized budgeted allocation of an anticipated marginal Lyapunov benefit instead of independent per-agent error thresholds, under a common heterogeneous plant, controller, communication graph, disturbance, and packet-loss model. It was observed that communication reduction is 92.2% under the new scheme compared to the periodic communication scheme and about 50% compared to the dynamic error-triggered communication scheme, and statistically indistinguishable regulation performance compared to the latter, implying that the new scheme, namely, to rank the candidate transmissions by the value of stability benefit associated with the communication cost of that transmission, is at least as efficient as reacting to an independently determined threshold in the local measurement. There are, however, several extensions needed to transform the numerical evidence from this simulation study into a formal contribution, fit for publication:

- Compute an upper bound for $S_{i,k}$ in closed-form or in a computable form that doesn't require evaluating the entire aggregate Lyapunov function at each candidate.
- Prove that the top-B allocator will maintain a desiredly prescribed Lyapunov margin (that is, formally show how a sufficient condition on B_k will ensure that $\varepsilon_{\text{comm}}(\gamma_k) \leq \varepsilon_{\max}$ in the ISS bound of Section 4).
- Include communication delay, quantization, and correlated (Markovian) packet loss directly in the analysis, as opposed to assuming that the packet loss is memoryless Bernoulli.
- To prevent Zeno-type behaviour in continuous-time extensions, periodically check the triggering/allocation condition.
- Evaluate the scalability for $N = 20{,}50{,}100$ agents to show the relationship between the amount of allocation cost ($\mathcal{O}(N \log N)$) and the amount of communication reduction obtained with an increase in the number of agents.
- Extend the Monte Carlo simulation from 8 to 30–50, or even more, different random initializations and report their confidence intervals; and compare with stability-constrained distributed MPC with event-triggered communication (Qin *et al*., 2025) directly.
- Perform controller-hardware-in-the-loop and power-hardware-in-the-loop validations, and follow the validation pathways for layers: exponentially triggered secondary control (Dehkordi*et al*., 2017; Lou *et al*., 2017; Al-Tameemi*et al*., 2026), which is the microgrid application described in Section 7.

This program, consisting of the following stages: ISS/ultimate bound explicit theorem, sufficient budget condition, and hardware validation, shall be the crucial and key step for distinguishing the current contribution from one based purely on simulation studies and is the major ongoing effort.

9. Conclusion

This paper presented a resource-allocation problem acting as a constraint on the control problem where, by means of a predicted marginal Lyapunov-stability benefit, communication scheduling in distributed networked control can be formulated. The proposed stability-budg*et al*locator does not ask whether a specific agent's state error goes beyond some threshold; it asks: which of the candidate transmissions sorts best among those currently eligible for allocation regarding the predicted stability gain per unit communication resource, and it answers that question with an O(N log N) ranking rule that is suitable for real-time applications. The proposed approach was tested via an eight-seed Monte Carlo study under a common heterogeneous test bed composed of a mix of lossy communication, stale information, bounded disturbances, and actuator saturation, and was shown to be able to reduce communication by 92.2% compared to periodic transmission and by approximately 50% compared to dynamic error-triggered communication without compromising baseline regulation performance, while tracing a well-behaved stability–communication Pareto frontier as the tight budget is swept. A distributed secondary voltage/frequency (VFC) control approach was presented for a direct application of inverter-based distributed energy resources (DERs) in islanded microgrids and a layered controller-HIL / power-HIL / testbed validation workflow. The major tasks are still theoretical work: a closed-form bound on the marginal stability benefit, a formal delay-quantization-packet-loss stability proof of the budgeted allocator, and finally, hardware-in-the-loop tests on a physical microgrid platform.

References

1. Al-Tameemi, Z. H. A., Peykarporsan, R., Lie, T. T., Zamora, R., & Blaabjerg, F. (2026). A comprehensive review of event-triggered consensus schemes in DC microgrids. Energies, 19(13), 2958. https://doi.org/10.3390/en19132958

2. Calpbinici, A., Irmak, E., &Kabalcı, E. (2024). Design and implementation of an energy management system with event-triggered distributed secondary control in DC microgrids. Energies, 17(3), 662. https://doi.org/10.3390/en17030662

3. Chen, M., Xiao, X., & Guerrero, J. M. (2018). Secondary restoration control of islanded microgrids with a decentralized event-triggered strategy. IEEE Transactions on Industrial Informatics, 14(9), 3870–3880. https://doi.org/10.1109/TII.2017.2784561

4. Dehkordi, N. M., Sadati, N., & Hamzeh, M. (2017). Distributed robust finite-time secondary voltage and frequency control of islanded microgrids. IEEE Transactions on Power Systems, 32(5), 3648–3659. https://doi.org/10.1109/TPWRS.2016.2634085

5. Ding, L., Han, Q.-L., Ge, X., & Zhang, X.-M. (2018). An overview of recent advances in event-triggered consensus of multiagent systems. IEEE Transactions on Cybernetics, 48(4), 1110–1123. https://doi.org/10.1109/TCYB.2017.2771560

6. Fan, Y., Feng, G., Wang, Y., & Song, C. (2013). Distributed event-triggered control of multi-agent systems with combinational measurements. Automatica, 49(2), 671–675. https://doi.org/10.1016/j.automatica.2012.11.010

7. Garcia, E., Cao, Y., & Casbeer, D. W. (2017). Periodic event-triggered synchronization of linear multi-agent systems with communication delays. IEEE Transactions on Automatic Control, 62(1), 366–371. https://doi.org/10.1109/TAC.2016.2555484

8. Ge, X., Yang, F., & Han, Q.-L. (2017). Distributed networked control systems: A brief overview. Information Sciences, 380, 117–131. https://doi.org/10.1016/j.ins.2015.07.047

9. Heemels, W. P. M. H., Teel, A. R., van de Wouw, N., &Nešić, D. (2010). Networked control systems with communication constraints: Tradeoffs between transmission intervals, delays and performance. IEEE Transactions on Automatic Control, 55(8), 1781–1796. https://doi.org/10.1109/TAC.2010.2042352

10. Hu, S., Zhang, Y., & Du, Z. (2012). Network-based H∞ tracking control with event-triggering sampling scheme. IET Control Theory & Applications, 6(4), 533–544. https://doi.org/10.1049/iet-cta.2011.0055

11. Huang, B., Xiao, Y., Jin, X., Feng, J., Li, X., & Ding, L. (2022). A novel dynamic event-triggered mechanism for distributed secondary control in islanded AC microgrids. Energies, 15(19), 6883. https://doi.org/10.3390/en15196883

12. Lou, G., Gu, W., Xu, Y., Cheng, M., & Liu, W. (2017). Distributed MPC-based secondary voltage control scheme for autonomous droop-controlled microgrids. IEEE Transactions on Sustainable Energy, 8(2), 792–804. https://doi.org/10.1109/TSTE.2016.2620283

13. Nešić, D., & Teel, A. R. (2004). Input–output stability properties of networked control systems. IEEE Transactions on Automatic Control, 49(10), 1650–1667. https://doi.org/10.1109/TAC.2004.835360

14. Olfati-Saber, R., & Murray, R. M. (2004). Consensus problems in networks of agents with switching topology and time delays. IEEE Transactions on Automatic Control, 49(9), 1520–1533. https://doi.org/10.1109/TAC.2004.834113

15. Olfati-Saber, R., Fax, J. A., & Murray, R. M. (2007). Consensus and cooperation in networked multi-agent systems. Proceedings of the IEEE, 95(1), 215–233. https://doi.org/10.1109/JPROC.2006.887293

16. Qin, D., Ong, C.-J., Liu, A., Zhang, W.-A., & Yu, L. (2025). Event-triggered distributed predictive control for multi-agent systems with stability constraints. Automatica, 178, Article 112358. https://doi.org/10.1016/j.automatica.2025.112358

17. Tabuada, P. (2007). Event-triggered real-time scheduling of stabilizing control tasks. IEEE Transactions on Automatic Control, 52(9), 1680–1685. https://doi.org/10.1109/TAC.2007.904277

18. Yang, F., & Han, Q.-L. (2013). H∞ control for networked systems with multiple packet dropouts. Information Sciences, 252, 106–117. https://doi.org/10.1016/j.ins.2013.06.043

19. Zhang, W., Bae, J., &Tomizuka, M. (2015). Modified preview control for a wireless tracking control system with packet loss. IEEE/ASME Transactions on Mechatronics, 20(1), 299–307. https://doi.org/10.1109/TMECH.2013.2297151

20. Zhang, X.-M., Han, Q.-L., Ge, X., Ding, D., Ding, L., Yue, D., & Peng, C. (2020). Networked control systems: A survey of trends and techniques. IEEE/CAA Journal of Automatica Sinica, 7(1), 1–17. https://doi.org/10.1109/JAS.2019.1911651