Selecting the Minimal Multi-Hop Radius for Resilient Consensus: A Hybrid Robustness–Proxy Framework for MW–MSR

Mohamed A. Sharaf *

Posted Date: 26 February 2026

doi: 10.20944/preprints202602.1237.v1

Keywords: resilient consensus; multi-agent systems; W–MSR; multi-hop communication; graph robustness; algebraic connectivity; power system test cases; adversarial agents

Article

Selecting the Minimal Multi-Hop Radius for Resilient Consensus: A Hybrid Robustness–Proxy Framework for MW–MSR

Mohamed A. Sharaf

Department of Computer Engineering and Networks, College of Information Sciences, Jouf University, Sakaka 72388, Saudi Arabia; masharaf@ju.edu.sa

Abstract

Achieving resilient consensus in adversarial environments often requires extending the W–MSR algorithm to multi-hop communication. While the robustness guarantees of multi-hop W–MSR are now well understood, the problem of how to determine the minimal hop radius h^* that ensures these guarantees has remained largely unaddressed. Existing work typically assumes a fixed h, leaving practitioners without a systematic way to balance robustness requirements against communication and computational cost. This paper introduces a new hop-selection framework that identifies the smallest communication horizon capable of satisfying the robustness assumptions underlying MW–MSR consensus. The framework combines exact robustness verification—when tractable—with a hierarchy of computationally efficient proxy tests based on local feasibility, normalized algebraic connectivity, and adversary-dilution criteria. These components provide a practical and scalable mechanism for establishing h^* in both synchronous and bounded-delay asynchronous settings. Design-time and runtime procedures, complexity analysis, and validation on IEEE 14-, 30-, and 57-bus networks demonstrate that the proposed approach reliably detects resilience thresholds and substantially improves consensus behavior under stealthy and burst-type adversaries. The results show that systematic hop selection is essential for avoiding failure at small h while preventing unnecessary communication overhead at large h. The framework thus offers an implementable and deployment-oriented strategy for resilient distributed coordination in sparse and adversarial multi-agent networks.

Keywords: resilient consensus; multi-agent systems; W–MSR; multi-hop communication; graph robustness; algebraic connectivity; power-system test cases; adversarial agents

1. Introduction

Achieving resilient consensus in distributed multi-agent systems is essential for ensuring stability and coordination in cyber–physical infrastructures, power networks, and networked control systems. The W–MSR (Weighted–Mean Subsequence Reduced) algorithm and its variants provide a foundational mechanism for tolerating Byzantine agents by filtering extreme neighbor values and maintaining convex-hull safety [1]. However, classical W–MSR relies on one-hop communication, which requires strong robustness properties (e.g., (r, s)-robustness) that many sparse networks—such as power-system graphs—cannot satisfy. To address this, recent extensions have introduced multi-hop W–MSR (MW–MSR) frameworks, where agents gather information from nodes up to h hops away [2]. These works establish robustness criteria for synchronous and asynchronous settings and show that multi-hop communication can significantly relax topological requirements. Specifically, if the h-hop graph satisfies $(f + 1)$-robustness (or $(2f + 1)$-robustness under delays), resilient consensus becomes achievable even when one-hop communication fails. Despite these advances, a critical practical challenge remains unresolved:

How can one determine the smallest hop radius h^* that guarantees resilient consensus while avoiding unnecessary communication overhead?

 doi:10.20944/preprints202602.1237.v1

In nearly all MW–MSR literature, the hop count h is preselected, often heuristically or as a fixed parameter. This leaves practitioners without guidance, forcing them to choose between either too small h that leads to insufficient robustness and consensus failure or too large h that leads to excess bandwidth usage, latency, and increased attack surface

Moreover, exact robustness verification—although theoretically ideal—is computationally difficult for large graphs and is often impractical in real deployments. As a result, there is a clear need for a systematic, scalable, and implementation-oriented procedure for selecting h^*.

Contributions– This paper introduces a new hop-selection framework that addresses this gap and is designed for practical deployment in large-scale, possibly sparse, adversarial networks. The key contributions are as follows: A hybrid hop-selection mechanism that integrates exact robustness verification (when tractable) with a set of computationally efficient proxy tests, enabling systematic identification of the minimal hop radius h^*.

A hierarchy of implementable proxies, including: local feasibility (minimum degree condition), normalized algebraic connectivity of the h-hop Laplacian, adversary-dilution constraints (when adversary locations are partly known). These proxies offer scalable approximations to robustness conditions while retaining practical effectiveness.

Design-time and runtime algorithms that are compatible with MW–MSR consensus under synchronous and bounded-delay asynchronous settings, ensuring consistency with the established theoretical guarantees.

Validation on IEEE 14-, 30-, and 57-bus network topologies, showing that the proposed framework reliably detects robustness thresholds, prevents failure under inadequate h, and avoids unnecessary communication overhead at excessive h.

Communication, computation, and latency analysis, quantifying trade-offs associated with hop expansion and supporting the framework's deployment in resource-constrained systems.

Overall, this work aims to provide a practical, scalable, and theoretically grounded approach to choosing the communication horizon for resilient consensus. Rather than proposing a new variant of W–MSR, our goal is to bridge the gap between MW–MSR theory and real-world implementation—offering tools that enable practitioners to configure multi-hop communication in a principled manner.

2. Related Work

Achieving consensus in multi-agent networks [3,4] that suffer dynamic updates or changes in the form of link or node failure, and delays is covered in Olfati et al.'s work [5]. The work in [5] is governed by sound theoretical, graphical and control theory.

To enable agents to come to an agreement in an environment that suffers malicious or faulty nodes is an area in which several algoritms based on W-MSR family are presented. The aim is to help normal agents to achieve consensus, see [1,6–8] for more details.

The presence of stealthy Byzantine agents with sufficient knowledge about the network and maliciously working to drift the system away from its normal behavior has been presented in Ishii's overview, see [9] and Zhao et al.'s work to isolate attacks [10].

The problem of resilient consensus in multi-agent leaderless systems with coordination is presented in [11,12].

The problem of resilient consensus in multi-agent in the presence of leaders has been studied greatly in literature. The work of [13] focuses on W-MSR algorithm to deals with time-varying grpahs unlike many models that deal with static graphs. However, Usevitch and Dimitra's work is limited to discrete-time dynamics with reliance on the W-MSR algorithm to achieve local filtering, see [14]. In addition, the work commonly counts on strong graph robustness ($r-$robust). Moreover, the proposed model is heavy in nature as it depends on complex leader dynamics.

Several efforts investigate the one-hop communication using W-MSR algorithm to reach almost sure consensus using matingale theory and random processes [15] and Shang's work [16]. Yemini

 doi:10.20944/preprints202602.1237.v1

et al. work relies on stochastic trust values [17]. Also, Rezaee et al. try to make mult-agent systems immune to DoS attack [18]. Rezaee et al. investigate W-MSR consensus under the presumption of noisy channels [19].

The notion of robustness and connectivity in complex networks are investigated in Zhang et al.'s work [20]. Also, Tyra et al.'s work [21] investigate models' robustness under adaptive and/or dependent attacks. The work [21] considers the attack scenarios that follow dependent and adaptive patterns.

Moreover, Usevitch and Dimitra's work is considered as the anchor for resilient leader-follower. The work by [22] is inspired by Usevitch and Dimitra's. However, Yuan and Ishii's work covered several gaps in [13] such as extending W-MSR by multi-hop relays, see [2,23] for more details on muti-hop communication.

Also, the work by Shang [24] provides a unified leaderless and leader-follower resilient consensus over directed random networks with l-hop communication, Byzantine nodes, and edge failure. This is different from this work in that Shang's work [24] introduces an l-hop communication as a generalization without claiming or proving "optimality". So, the distinction between this work and Shang's [24] is that Shang is after a feasible framework while this work is targeting (performance-driven) optimal design. Also, this work seeks optimal hop radius h^* by merging a proven heuristic and exact search, and validates on IEEE-14/30/57 bus systems.

According to Abbas et al. [25], the presence of trusted nodes is a game changer and helps improve robustness and resilient consesus.

Also, the work by Niewenhuis and Varbanescu [**?**] addresses trimming concept based on the well-established priniciple of strongly connected components (SCCs), see [26]. Niewenhuis and Varbanescu introduce a novel algorithm "Forward-Backward" (FB)to compute SCCs. The importance of Trimming in M-MSR algorithm is due to the needto eliminate adversaries.

3. Preliminaries

Let $G = (\mathcal{V}, \mathcal{E})$ be a directed or undirected communication graph on $|\mathcal{V}| = N$ agents. The adjacency matric is denoted by A. For a positive integer h, define the h-hop adjacency matrix as in Equation (1):

$$A^{(h)} = sgn(\sum_{k=1}^{h} A^k), \tag{1}$$

with a zero diagonal. The corresponding h-hop neighborhood of agent i is

$$N_i^{(h)} = \{j \in V : A_{ij}^{(h)} = 1\},$$

and the h-hop degree is $deg_i^{(h)} = |N_i^{(h)}|$.

3.1. Robustness Concepts

For resilience to Byzantine adversaries, W–MSR and MW–MSR rely on the well-established framework of (r, s)-robustness. A graph is (r, s)-robust if every pair of nonempty, disjoint subsets of nodes contains at least one subset where at least s nodes have at least r neighbors outside the subset.

For a given hop radius h, define the h-hop robustness index of G as shown in Equation (2):

$$R_h(G) = max\Big\{k : G^{(h)} is(k, f+1) - robust\Big\}. \tag{2}$$

Theorem 1 (Synchronous MW-MSR [2]). *Consensus is achieved iff the h-hop graph $G^{(h)}$ is $(f+1, f+1)$-robust.*

Theorem 2 (Asynchronous with bounded delays [2]). *$(f+1, f+1)$-robustness is necessary and $(2f+1)$-robustness is sufficient.*

 doi:10.20944/preprints202602.1237.v1

These results directly motivate the search for a minimal h that satisfies the corresponding robustness requirement.

3.2. Monotonicity of Multi-Hop Robustness

A key property of MW–MSR is that robustness does not decrease with additional hops:

$$h_1 \leq h_2 \rightarrow R_{h_1}(G) \leq R_{h_2}(G).$$

This monotonicity implies that if a certain robustness condition holds at some hop h_0, it will continue to hold for all $h \geq h_0$. This property is central to establishing minimality of the selected hop radius h^*.

4. Model and Adversary

Consider a network of N agents indexed by $V = \{1, \ldots, N\}$ and communicating over an underlying graph $G = (V, E)$. Each agent i evolves according to second-order dynamics:

$$p_i = v_i, \qquad \dot{v}_i = u_i,$$

where $p_i, v_i, u_i \in \mathcal{R}^d$ denote position, velocity, and control input, respectively. Agents exchange state information through their h-hop neighborhood $N_i^{(h)}$, determined by the h-hop adjacency matrix $A^{(h)}$ constructed in Section 3.

4.1. Normal and Adversarial Agents

The set of agents is partitioned as:

$$V = N \cup A, \quad N \cap A = \phi,$$

where N denotes normal agents that follow the prescribed multi-hop W-MSR protocol and A denotes adversarial agents.

We adopt an f-total adversarial model, where the total number of adversaries satisfies:

$$|A| \leq f_{max}.$$

Adversarial agents may deviate arbitrarily from the dynamics above and may send arbitrary, inconsistent, or malicious values to their neighbors. They may also coordinate or behave strategically based on global knowledge of the system.

4.2. Adversary Capabilities

The adversary model encompasses stealthy drift and burst-type attacks, both consistent with practical threat scenarios and prior work on Byzantine multi-agent systems. In particular, adversarial agents may:

1. send false state values (e.g., biased, extreme, or time-varying corrupted measurements).
2. Introduce intermittent bursts, temporarily pushing the normal agents away from consensus.
3. Exploit network sparsity, attempting to dominate local neighborhoods in low-degree regions.
4. Remain stealthy, keeping their transmitted values within plausible bounds to evade trimming early in the evolution.

These behaviors align with adversarial models used in recent resilient consensus studies and intentionally stress the W–MSR filtering mechanism.

 doi:10.20944/preprints202602.1237.v1

4.3. Communication Under Multi-Hop W-MSR and Trimming

At each update cycle, agent i collects pairs (p_j, v_j) from all agents in $N_i^{(h)}$, where information is relayed over up to h hops. The effective neighborhood depends on the hop radius, and therefore the robustness guarantees of the overall system directly depend on the chosen h. To ensure that the MW–MSR filter is well-posed, each agent i must be able to discard/trim up to $f_i^{(h)}$ potentially malicious values on each side of the coordinatewise ordering, yielding the feasibility condition:

$$deg_i^{(h)} \geq 2f_i^{(h)} + 1.$$

The value of $f_i^{(h)}$ is defined as in shown in Equation (3):

$$f_i^{(h)} = \min\left\{f_{\max}, \left\lfloor \frac{\deg_i^{(h)} - 1}{2} \right\rfloor\right\}. \tag{3}$$

This condition guarantees that, after trimming, at least one neighbor value remains available for computing safe averages.

5. Multi-Hop W-MSR Control Law

Under the MW–MSR framework, each normal agent collects multi-hop information, discards potentially malicious outliers, and applies a distributed control input computed from safely filtered neighbor data. The steps below formalize information gathering, trimming, and control.

5.1. Multi-Hop Information Gathering

Given a hop radius h and the associated adjacency $A^{(h)}$, agent i receives state pairs (p_j, v_j) from all $j \in N_i^{(h)}$. Messages may be relayed through intermediate nodes, so adversarial values can appear anywhere along the paths. This expanded neighborhood increases information redundancy but also enlarges the set of potentially adversarial inputs, hence the need of a properly sized trimming budget $f_i^{(h)}$.

5.2. Trimming Rule (WM-MSR)

At each update, agent i processes the received state values coordinatewise:

1. Sort the collected values of p_j and v_j for all $j \in N_i^{(h)}$.
2. Trim the largest $f_i^{(h)}$ and smallest $f_i^{(h)}$ values.
3. Retain the remainder; feasibility is ensured by

$$\deg_i^{(h)} \geq 2f_i^{(h)} + 1, \qquad f_i^{(h)} = \min\left\{f_{\max}, \left\lfloor \frac{\deg_i^{(h)} - 1}{2} \right\rfloor\right\}$$

Let $\mathcal{R}_i^{(h)}$ denote the retained set of neighbors after trimming. The filtered averages are

$$\bar{p}_i = \frac{1}{|\mathcal{R}_i^{(h)}|} \sum_{j \in \mathcal{R}_i^{(h)}} p_j, \quad \bar{v}_i = \frac{1}{|\mathcal{R}_i^{(h)}|} \sum_{j \in \mathcal{R}_i^{(h)}} v_j,$$

consistent with the WM-MSR message-cover interpretation.

5.3. Distributed Control Input

The control law of agent i is defined as:

$$u_i = -\alpha(p_i - \bar{p}_i) - \beta(v_i - \bar{v}_i), \quad \alpha, \beta > 0,$$

 doi:10.20944/preprints202602.1237.v1

i.e., a Proportional-Derivative (PD)-type consensus controller that drives each agent toward filtered multi-hop references. Given the trimming feasibility and the robustness conditions on $G^{(h)}$, this preserve convex-hull safety of the normal agents and ensure resilient convergence under the appropriate model (synchronous or bounded-delay asynchronous).

5.4. Continuous/Discrete Operation

At each operation: (i) gather(p_j, v_j) from $N_i^{(h)}$; (ii) trim $\pm f_i^{(h)}$ extremes; (iii) compute $(\bar{p}_i, \bar{v}_i)$; (iv) update u_i. The loop repeats until completion. When the robustness requirement at the chosen h holds, MW-MSR achieves resilient asymptotic consensus in the corresponding model.

6. Optimal Hop Selection

Figure 1 shows a conceptual overview of the proposed hop-selection framework.

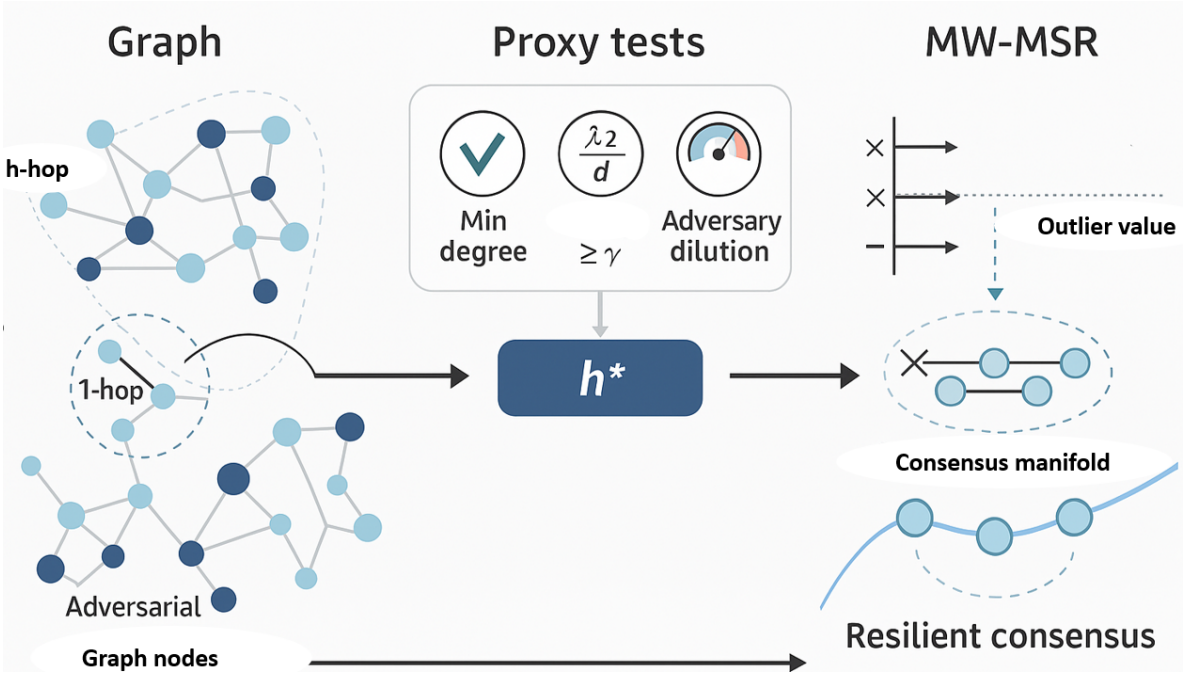

Figure 1. Conceptual diagram of the optimal multi-hop W-MSR framework.

Choosing h is central: too small—insufficient robustness; too large—unnecessary bandwidth, latency, and attack surface. We formalize the optimal hop radius and present a hybrid selection procedure that identifies the smallest feasible h^* using exact robustness verification (when available) combined with lightweight proxies.

6.1. Robustness Targets and h^*

Let

$$P_{\text{syn}}(h) \equiv G^{(h)} \text{ is } (f+1, f+1)\text{-robust}, \quad P_{\text{asyn}}(h) \equiv G^{(h)} \text{ is } (2f+1)\text{-robust}.$$

Define

$$h^*_{\text{syn}} = \min\{h \geq 1 : P_{\text{syn}}(h)\}, \quad h^*_{\text{asyn}} = \min\{h \geq 1 : P_{\text{asyn}}(h)\}.$$

These are the smallest horizons meeting the MW–MSR robustness assumptions for synchronous and bounded-delay asynchronous models, respectively.

6.2. Existence and Monotonicity

If $\mathcal{P}(h_0)$ holds for some h_0 then h^* exists with $h^* \leq h_0$, and $\mathcal{P}(h_0)$ holds for all $h \geq h^*$ due to the monotonicity of $h-$hop robustness with respect tp h.

 doi:10.20944/preprints202602.1237.v1

6.3. Principle of Optimality

For $h < h^*$, robustness fails and consensus cannot be guaranteed; at $h = h^*$, the robustness requirement holds; for $h > h^*$, correctness does not improve while costs strictly increase.

6.4. Hybrid Verification Strategy

Because exact (r, s)-robustness checking is expensive at scale, we proceed in two tiers:

- Tier-1 (Exact): If an exact checker or MILP is available and certifies $P(h)$, set $h^* = h$.
- Tier-2 (Proxies): Otherwise, apply the following scalable checks:

$$\min_i \deg_i^{(h)} \geq 2f_{max} + 1, \quad \frac{\lambda_2\left(L^{(h)}\right)}{\deg^{(h)}} \geq \gamma, \quad \max_i \frac{\left|N_i^{(h)} \cap A\right|}{\left|N_i^{(h)}\right|} \leq \rho_{\max},$$

capturing local feasibility, spectral connectivity, and adversary dilution (ρ_{max} denotes maximum allowed adversarial concentration).

6.5. Selection Procedure

Scan $h = 1, 2, \ldots, h_{max}$. If exact verification succeeds, return h^*. Otherwise, record the smallest h that passes all proxies; if none pass, return $\arg\max_h \lambda_2(L^{(h)})/\deg^{(h)}$ as a best-effort choice.

6.6. Cost Model and Trade-Offs

We minimize the cost expressed in Equation (4)

$$C_{\text{total}}(h) = \lambda_b C_{\text{bw}}(h) + \lambda_l C_{\text{lat}}(h) + \lambda_c C_{\text{comp}}(h), \tag{4}$$

where $\lambda_b + \lambda_l + \lambda_c = 1$ weight bandwidth ($\propto \deg^{(h)}$), latency ($\propto h$), and computation dominated by $O\left(\left|N_i^{(h)}\right| \log\left|N_i^{(h)}\right|\right)$. This underscores the desirability of smallest feasible h^*. Although we do not experimentally evaluate these costs, the qualitative model highlights why selecting the smallest feasible h^* is desirable.

6.7. Practical Interpretation

- $h < h^*$: Failure(drift/partition).
- $h = h^*$: Success with optimal efficiency.
- $h > h^*$: Success with redundant overhead.

This classification emphasizes why the smallest feasible hop radius is desirable for resilient and efficient operation.

6.8. Generalization to Directed Graphs (Digraphs)

All components of the framework extend to directed graphs by constructing $A^{(h)}$ and $L^{(h)}$ for the digraph and interpreting (r, s)-robustness in the directed sense (external in-neighbors). The monotonicity of h-hop reachability remains, so the principle of optimality and the hybrid selection logic are unchanged. The spectral proxy $\lambda_2(L^{(h)})$ should be interpreted via the symmetrized Laplacian or other digraph connectivity surrogates, while the local-feasibility and dilution proxies apply verbatim. Furthermore, for digraphs, one can use in-degree Laplacian ($L_{in} = D_{in} - A$) as a baseline, where D_{in} is the diagonal matrix of in-degrees and A is the adjacency matrix. Another variant for digraphs is row-stochastic weighted matrix.

7. Algorithmic Perspective of the Optimal Multi-Hop W-MSR Framework

Algorithm 1 exhibits the following characteristics: Graph-theoretic (since it is based on network connectivity properties), multi-objective optimization (as it balances horizon–communication cost–

 doi:10.20944/preprints202602.1237.v1

against robustness) and Hybrid search algorithm (since it combines exact verification with heuristic checks– as mentioned earlier).

Algorithm 1 Hop-Radius Selection with Robustness Verification

Input: Adjacency A, $f_{\max}$, $h_{\max}$, target $P \in \{\text{sync}, \text{async}\}$, thresholds γ, $\rho_{\max}$
Output: Selected hop radius h^*
$h^* \leftarrow$ null, $h_{\text{tent}} \leftarrow$ null
for $h = 1$ to $h_{\max}$ **do**
  Build $A^{(h)}$, $f^{(h)}$ via Algorithm 3
  if EXACTCHECKAVAILABLE() **then**
    if EXACTCHECK(P, h) **is true then**
      $h^* \leftarrow h$; **break** ▷ Exact minimal solution
    end if
  end if
  $pass \leftarrow$ SANITYCHECK$(A^{(h)}, f_{\max}, \gamma, \rho_{\max})$
  if $pass$ **and** $h_{\text{tent}} =$ null **then**
    $h_{\text{tent}} \leftarrow h$
  end if
end for
if $h^* =$ null **then**
  $h^* \leftarrow h_{\text{tent}}$
end if
return h^*

Next. Algorithm 2 implements the core resilient consensus controller that addresses the following: collects multi-hop state information, filters potentially malicious extreme values, computes resilient local coverage, applies consensus control flow and runs (loops) continuously untill stopped (running = false).

Algorithm 2 Agent-Level MW–MSR Controller (at agent i)

Input: $A^{(h)}$, $f_i^{(h)}$, gains $\alpha, \beta > 0$
while RUNNING **do**
  Receive (p_j, v_j) from all $j \in N_i^{(h)}$
  Trim top $f_i^{(h)}$ and bottom $f_i^{(h)}$ values (coordinatewise)
  Compute $\bar{p}_i, \bar{v}_i$ as averages of remainder
  $u_i \leftarrow -\alpha\,(p_i - \bar{p}_i) - \beta\,(v_i - \bar{v}_i)$
end while

Algorithm 3 provides the multi-hop consensus with the following: extends communication ranges from 1-hop to h-hop neighborhoods, quantifies resilience capacity (for each agent based on the extended connectivity), and enables the W-MSR to leverage multi-hop information for better fault tolerance. It is the key factor behind having the proposed optimal multi-hop W-MSR more tolerant to more adversarial agents while achieving consensus.

Algorithm 4 is responsible for systematically, adaptively selecting the communication horizon for resilient consensus by accomplishing the following: seeks minimality (reduces communication overhead), guarantees feasibility (always returns a solution), balances accuracy versus computation (follows hierarchical verification), and adapts and scale to problem size (smaller/denser graphs lead to smaller sufficient h^* while larger/sparser graphs necessitates larger h^*).

 doi:10.20944/preprints202602.1237.v1

Algorithm 3 h-Hop Adjacency and Trim-Budget Computation

Input: Binary adjacency A (zero diagonal), hop h, $f_{\max}$
Output: $A^{(h)}$, $f^{(h)}$
$A^{(h)} \leftarrow \mathbf{0}$; $B \leftarrow A$
for $k = 1$ to h **do**
  $A^{(h)} \leftarrow A^{(h)} \vee (B \neq \mathbf{0})$ ▷ Boolean OR of nonzeros
  $B \leftarrow \textsc{Sgn}(B\,A)$ ▷ Boolean sparse product
end for
Zero the diagonal of $A^{(h)}$; compute degrees $\deg_i^{(h)}$
for each node i **do**
  $f_i^{(h)} \leftarrow \min\left\{ f_{\max}, \left\lfloor \frac{\deg_i^{(h)} - 1}{2} \right\rfloor \right\}$
end for

Algorithm 4 Optimal Horizon Finder (Design-time)

Input: A, $f_{\max}$, $h_{\max}$, target P, thresholds $\gamma, \rho_{\max}$
Output: h^*
for $h = 1$ to $h_{\max}$ **do**
  Build $A^{(h)}$, $f^{(h)}$ via Algorithm 3
  if $\textsc{ExactCheck}(P, h)$ **true then return** $h^* = h$
  end if
  if $\textsc{SanityCheck}(A^{(h)}, f_{\max}, \gamma, \rho_{\max})$ **passes and** no smaller h passed **then**
    Tentatively set $h^* \leftarrow h$
  end if
end for
if h^* is defined **then return** h^*
elsereturn $\arg\max_h \ \lambda_2(L^{(h)}) / \deg^{(h)}$
end if

Algorithm 5 works as a screening tool to evaluate whether an h-hop could support resilient consensus. It accomplishes this task by checking: local connectivity (minimum degree, each node i needs at least $2f + 1$ neighbors), global connectivity (normalized algebraic connectivity $\lambda_2(L^{(h)})$), and adversarial distribution (dilution in neighborhoods). The sanity check algorithm systematically identifies the smallest h that passes its checks.

Algorithm 5 Sanity Check at Hop h

Input: $A^{(h)}$, $f_{\max}$, γ; (optional) adversary mask, $\rho_{\max}$
Output: PASS/FAIL
Compute $\deg_i^{(h)}$ for all i
if $\min_i \deg_i^{(h)} < 2f_{\max} + 1$ **then**
  return FAIL
end if
Compute $\lambda_2(L^{(h)})$ and $\deg^{(h)}$
if $\lambda_2(L^{(h)}) / \deg^{(h)} < \gamma$ **then**
  return FAIL
end if
if adversary mask given **then**
  if $\max_i |N_i^{(h)} \cap A| / |N_i^{(h)}| > \rho_{\max}$ **then**
    return FAIL
  end if
end if
return PASS

 doi:10.20944/preprints202602.1237.v1

8. Experimental Results

This section evaluates the proposed hop-selection framework and multi-hop W–MSR controller on IEEE 14-, 30-, and 57-bus power-network topologies under stealthy and burst-type adversarial behavior. All experiments compare single-hop W–MSR ($h = 1$) against multi-hop W–MSR with the selected horizon h^* obtained using the framework selection method proposed in Section 6. Table 1 shows simulation's parameters and assigned values.

Table 1. Simulation's parameters.

Parameter Synthesis	Value
Sample time, T_s	0.05 second
Maximum number of adversarial neighbors tolerable per node, f_{max}	2
Number of maximum hops, h_{max}	7
Gains (α, β)	2 and 3
Adversarial concentration (ρ_{max})	0.25
Number of adversaries (num_adv)	10% of the nodes in the network
Number of nodes, N	14, 30 and 57
Resilience threshold, γ	0.25

The results highlight three main outcomes:

1. Selecting h^* is essential for avoiding drift and ensuring resilient consensus.
2. The normalized algebraic connectivity $\lambda_2\left(L^{(h)}\right) / \deg^{(h)}$ effectively predicts the resilience threshold.
3. Increasing h beyond h^* offers no additional correctness benefit and only increase communication and computational load.

All experiments use up to 10% adversarial nodes, chosen uniformly at random unless otherwise noted.

8.1. Disagreement Norm Under Adversaries

Figure 2 compares the disagreement norm $\left\|x_i(t) - x_j(t)\right\|$ between normal nodes under two cases: single-hop W-MSR, and multi-hop W-MSR using the selected horizon h^*.

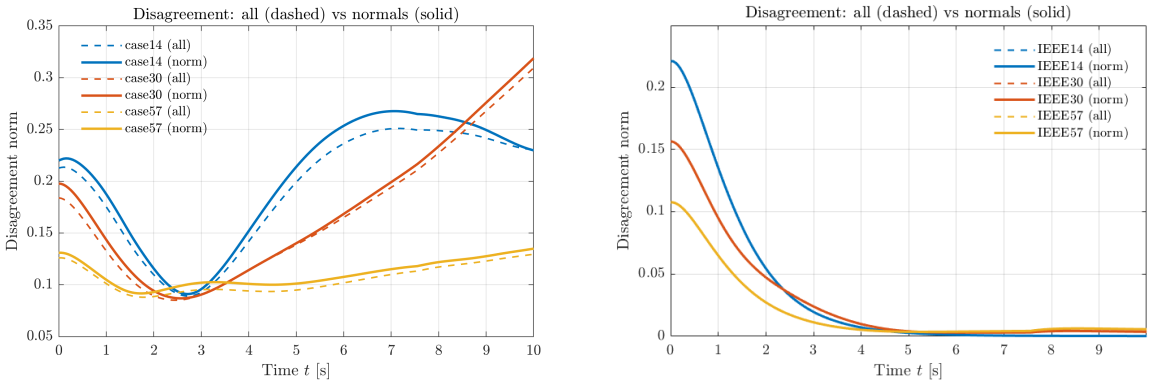

(**a**) (**b**)

Figure 2. Disagreement norm $\left\|x_i(t) - x_j(t)\right\|$ over time for IEEE 14-, 30-, and 57-bus systems. (a) Single-hop W–MSR ($h = 1$) fails to suppress adversarial drift, resulting in persistent disagreement. (b) Multi-hop W–MSR with the selected hop radius h^* achieves resilient consensus, with disagreement converging to zero for all normal agents.

Key observations

- Single-hop W–MSR fails to suppress adversarial drift (Figure 2 (a)), with disagreement remaining nonzero and often growing over time.

 doi:10.20944/preprints202602.1237.v1

- With the selected h^*, disagreement converges to zero, demonstrating resilient consensus across all IEEE graphs (Figure 2 (b)).
- For sparse networks like IEEE 57-bus, hop augmentation is critical; consensus is impossible when agents rely only on one-hop information.

This validates that h^* restores sufficient robustness where single-hop communication is inadequate.

8.2. Final-State Spread (SSE) and Effect of Connectivity

Final convergence accuracy is measured using the steady-state error (SSE): error using all nodes (SSE-all) and error using only normal nodes (SSE-normal).

Figure 3 (a) and (b) illustrate how the hop radius affects SSE.

Key findings

- At $h = 1$, adversaries dominate local neighborhoods in sparse regions, producing large SSE values (e.g., SSE ≈ 2.71 for IEEE-57).
- At the selected h^*, SSE decreases dramatically (e.g., SSE ≈ 0.07), and SSE-all nearly matches SSE-normal, indicating that adversaries can no longer partition or significantly bias the network.
- This improvement correlates with an increase in the normalized algebraic connectivity $\lambda_2(L^{(h)})$, confirming its usefulness as a resilience proxy.

Overall, h^* provides the required "global visibility" that defeats influence concentration by adversaries.

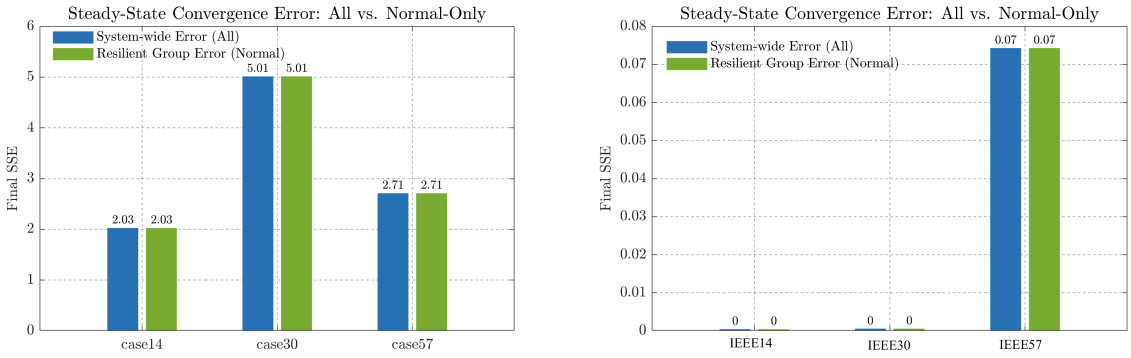

(**a**) (**b**)

Figure 3. Final steady-state error (SSE) for all agents ("SSE-all") and normal agents only ("SSE-normal"). (a) At $h = 1$, adversaries dominate sparsely connected neighborhoods, resulting in large SSE values (e.g., SSE ≈ 2.71 for IEEE-57). (b) At the selected hop radius h^*, SSE is significantly reduced (e.g., SSE ≈ 0.07 for IEEE-57), and SSE-all nearly matches SSE-normal, indicating effective suppression of adversarial influence..

8.3. Control-Energy Profile

Control efficiency is examined using the instantaneous control energy $\|u_i(t)\|$. Figure 4 illustrates that in the single-hop case, the control energy exhibits frequent bursts due to persistent adversarial disturbances. Larger networks show more severe spikes. With optimal multi-hop, control energy rapidly decays to zero once consensus is reached, indicating efficient and stable convergence.

This distinction offers a practical diagnostic: If both disagreement and control energy remain persistently nonzero, the chosen hop radius is insufficient.

 doi:10.20944/preprints202602.1237.v1

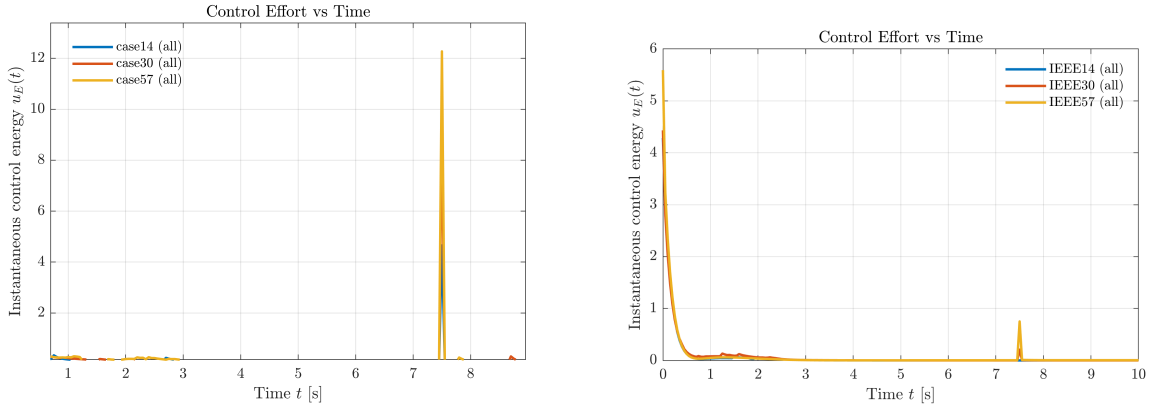

(**a**) (**b**)

Figure 4. Instantaneous control energy $\|u_i(t)\|$ for normal agents. (a) Under single-hop W–MSR, adversarial disturbances lead to repeated spikes in control effort, with larger networks showing more pronounced bursts. (b) At the selected hop radius h^*, control energy rapidly decays as consensus is established, demonstrating efficient and stable convergence.

8.4. Convex-Hull Evolution

Figure 5 examines the evolution of the convex hull of the normal agents' states.

Observations

- For $h = 1$, the convex-hull envelope expands over time, reflecting instability and adversarial influence.
- At h^*, hull expansion becomes much flatter and tightly bounded, even for the IEEE-57 graph, demonstrating robust containment of adversarial drift.

This agrees with the MW–MSR safety property that normal states remain within the convex hull of initial normal values when the robustness condition is satisfied.

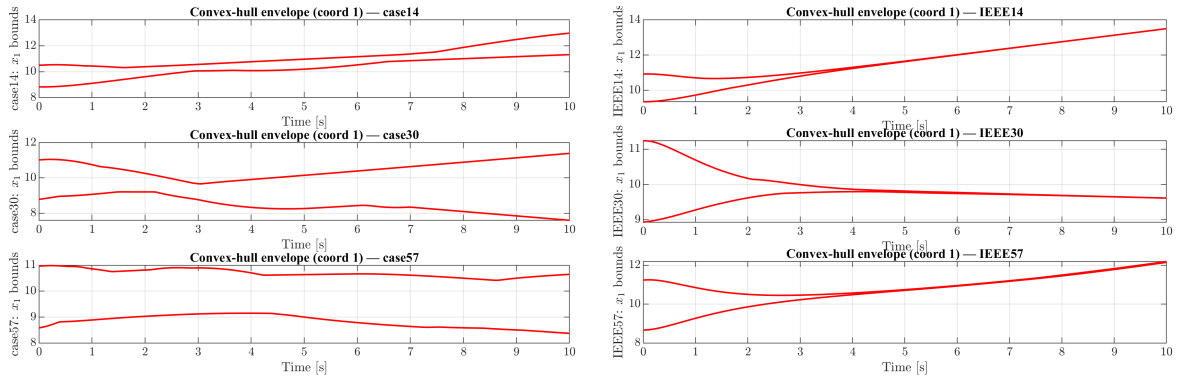

(**a**) (**b**)

Figure 5. Evolution of the convex hull of normal agents' states. (a) For $h = 1$, the convex-hull envelope grows over time, reflecting the inability of single-hop filtering to contain adversarial drift. (b) At h^*, the hull remains tightly bounded, even for sparse networks such as IEEE-57, demonstrating robust containment of adversarial influence.

8.5. Scalability and Behavior Across IEEE Graphs

The proposed hop-selection method successfully adapts to the structural differences between IEEE-14, IEEE-30, and IEEE-57:

- Dense subgraphs (e.g., IEEE-14) require smaller horizons.
- Sparse or large-diameter networks (e.g., IEEE-57) require larger horizons to achieve adequate robustness.

 doi:10.20944/preprints202602.1237.v1

- Across all cases, the method identifies an h^* that prevents failure modes at insufficient hop counts while avoiding unnecessary overhead at larger values.
- For IEEE-57, the selection $h^* = 6$ reorganizes the sparse graph into a sufficiently connected multi-hop structure to support consensus.

This demonstrates that the proposed framework scales naturally to larger and more challenging networks.

8.6. Sequence Diagram of the Proposed Framework

Figure 6 provides a sequence diagram summarizing the flow from hop-selection to MW–MSR execution. The diagram follows these steps:

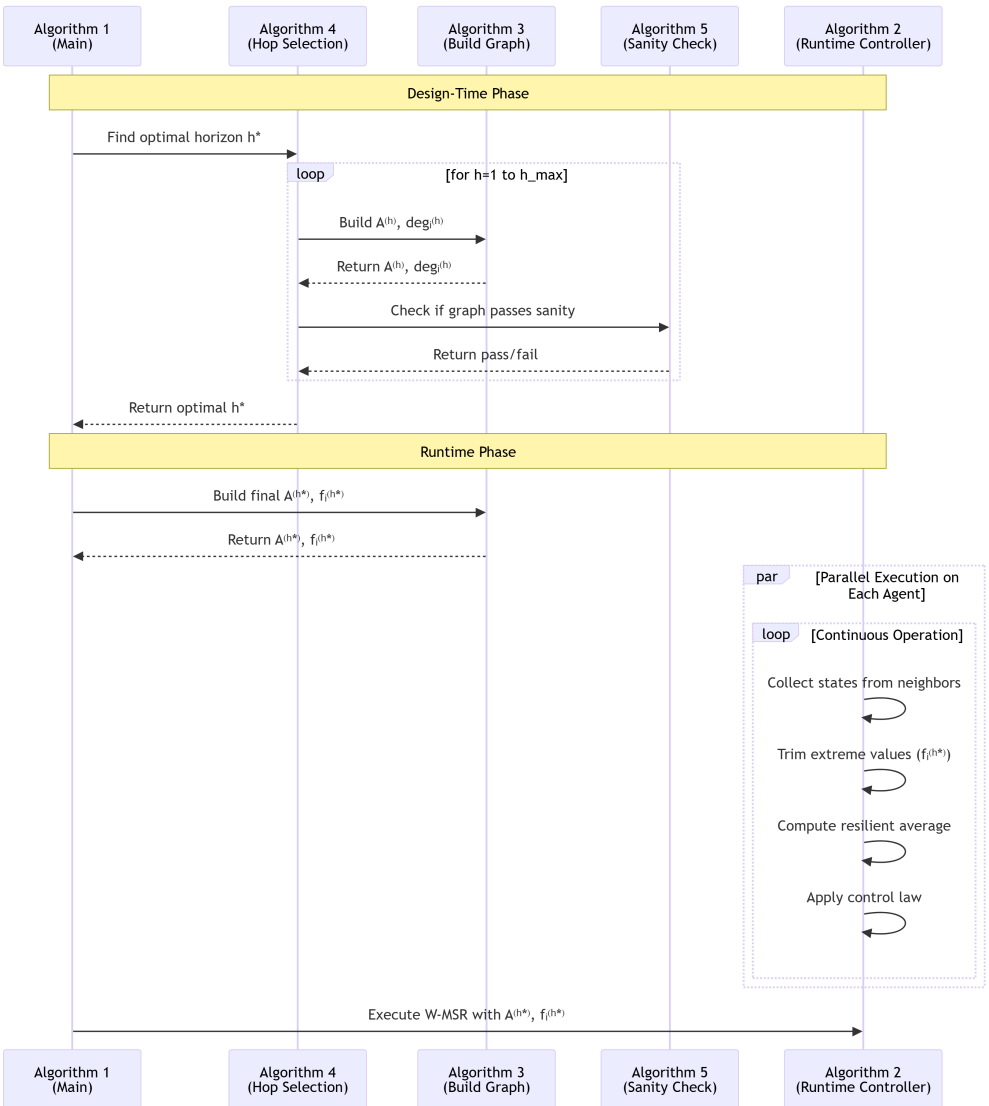

Figure 6. Sequence diagram of the proposed optimal multi-hop W-MSR model.

1. Construct h-hop adjacency.
2. Apply exact or proxy check.
3. Identify the minimal hop radius h^*.
4. Execute MW-MSR with the corresponding per-node trim budgets.

This offers a clear visualization of how the proposed hybrid method integrates design-time and runtime components.

 doi:10.20944/preprints202602.1237.v1

9. Discussion

The proposed hop-selection framework provides a structured and practical approach to determining the minimal communication horizon required for resilient consensus under the MW–MSR algorithm. The results presented in Section 7 demonstrate that resilience depends not only on the network topology but also on the selection of an appropriate hop radius that compensates for sparsity and adversarial influence.

Adaptation to Network Topology

The experiments highlight that each IEEE network exhibits distinct structural properties that influence the required hop radius. For example, the IEEE-57 bus system is designed for efficient power distribution, not for consensus or information fusion. Its sparsity and relatively large diameter make it unsuitable for single-hop W–MSR, leading to persistent drift and disagreement. The hop-selection framework automatically identifies a sufficiently large horizon—h^* in the experiments—to overcome these structural limitations and enable resilient consensus.

This adaptive behavior is important because it demonstrates that the method does not rely on manual tuning or overly conservative hop choices. Instead, it systematically adjusts the communication radius to meet resilience requirements dictated by the graph structure.

Effectiveness Against Stealthy and Burst Adversaries

The proposed framework is tested against adversaries capable of both stealthy drift and burst-type disturbances. These adversaries attempt to exploit local sparsity or low-degree regions to bias normal agents or cause partial divergence. The results show that:

- When $h < h^*$, adversaries can dominate local neighborhoods, leading to high SSE and unstable convex-hull behavior.
- When $h = h^*$, adversarial influence is effectively diluted, neighborhood redundancy increases, and the filtered averages remain reliable.

The alignment of SSE-normal and SSE-all at h^* confirms that the adversaries can no longer distort the global consensus trajectory.

Role of Algebraic Connectivity

The experiments further validate the role of normalized algebraic connectivity $\lambda_2(L^{(h)})/\deg^{(h)}$ as a useful resilience indicator. Increases in λ_2 correlate with improved consensus performance and reduced vulnerability, particularly in sparse networks. Although λ_2 is not a substitute for exact robustness checking, it provides a computationally efficient and reliable proxy that integrates naturally into the hop-selection pipeline.

Efficiency and Resource Awareness

The results reinforce the importance of selecting the smallest feasible hop radius. Larger hop counts expand the communication graph but increase the number of relayed messages, latency, and computational workload. The proposed framework avoids these unnecessary costs by prioritizing minimality. For example, while increasing h beyond h^* does not harm consensus correctness, it leads to superfluous overhead without further resilience benefits.

Scalability and Practical Deployment

The hop-selection framework provides a scalable strategy for determining communication horizons in large real-world networks. The combination of exact verification (when feasible) and efficient proxy tests (when exact checking is impractical) ensures that the algorithm can operate effectively across networks of varying sizes and densities. The use of multi-hop communication, combined with adaptive hop selection, enables the MW–MSR controller to operate reliably even in networks with challenging topology.

 doi:10.20944/preprints202602.1237.v1

Additionally, because the framework requires only adjacency information and standard graph computations, it can be deployed in settings where computational resources are limited, making it suitable for cyber–physical systems, power networks, distributed robotics, and IoT applications.

10. Conclusion and Future Work

This paper introduces a structured framework for selecting the minimal hop radius required for resilient consensus under the multi-hop W–MSR (MW–MSR) algorithm. The proposed approach integrates exact robustness verification—when computationally feasible—with a set of lightweight and scalable proxy tests involving local feasibility, normalized algebraic connectivity, and adversary-dilution metrics. These components together enable a principled and practical mechanism for identifying the smallest communication horizon h^* that satisfies the robustness assumptions of MW–MSR in both synchronous and bounded-delay asynchronous settings.

Experimental results on IEEE 14-, 30-, and 57-bus systems confirm that selecting h^* is essential for resilient operation. When $h < h^*$, adversaries are able to exploit sparsity, induce drift, or expand the convex hull of normal states, resulting in consensus failure. In contrast, at $h = h^*$, the multi-hop neighborhoods provide sufficient structural redundancy to suppress adversarial influence, yielding significantly improved disagreement, SSE performance, and control-energy behavior. The method scales naturally with network size and topology, and it avoids the communication and computation overhead associated with unnecessarily large hop values.

Looking ahead, several research directions can extend the usefulness of the proposed framework. First, developing more scalable exact robustness certification techniques—for example, through mixed-integer formulations or convex relaxations—would improve accuracy for large networks. Second, incorporating cost-aware hop selection based on latency, bandwidth, and energy budgets may enable deployment in resource-constrained settings. Third, extending the method to time-varying graphs using windowed or dynamic h-hop robustness measures is a natural next step. Finally, adaptive mechanisms that adjust $h(t)$ in real time, as well as extensions to vector-valued or privacy-preserving consensus, offer promising directions for future exploration.

Conflicts of Interest: Declare conflicts of interest or state "The authors declare no conflict of interest." Authors must identify and declare any personal circumstances or interest that may be perceived as inappropriately influencing the representation or interpretation of reported research results. Any role of the funders in the design of the study; in the collection, analyses or interpretation of data; in the writing of the manuscript; or in the decision to publish the results must be declared in this section. If there is no role, please state "The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results".

Appendix A. Local-Feasibility Lemma and Complexity

Lemma A1 (Local feasibility: trim budget vs degree)**.** *Let* $\deg_i^{(h)}$ *be the number of neighbors of agent* i *in the* h*-hop graph* $G^{(h)}$*, excluding self-loops. Suppose that*

$$\deg_i^{(h)} \geq 2f_i^{(h)} + 1,$$

where

$$f_i^{(h)} = \min\{f_{\max}, \lfloor (\deg_i^{(h)} - 1)/2 \rfloor\}.$$

Then, after discarding the $f_i^{(h)}$ *largest and* $f_i^{(h)}$ *smallest neighbor values (coordinatewise), the remaining set is nonempty. Consequently, the MW–MSR update for agent* i *is well-defined.*

Proof. Agent i has $\deg_i^{(h)}$ neighbors in $G^{(h)}$. Trimming removes at most $2f_i^{(h)}$ values. By assumption,

$$\deg_i^{(h)} - 2f_i^{(h)} \geq 1,$$

 doi:10.20944/preprints202602.1237.v1

so at least one neighbor value remains after trimming. Therefore, the filtered set of values used to compute the MW–MSR update is nonempty, and the average used to compute $(\bar{p}_i, \bar{v}_i)$ is well-posed.

Because

$$f_i^{(h)} = \min\left\{ f_{\max}, \left\lfloor \frac{\deg_i^{(h)} - 1}{2} \right\rfloor \right\}$$

is constructed to satisfy

$$2f_i^{(h)} + 1 \le \deg_i^{(h)},$$

the feasibility condition always holds. □

Appendix A.1. Complexity of the Hop-Selection Pipeline

The complexity of the hop-selection framework arises primarily from constructing the multi-hop adjacency matrices, computing spectral quantities, and performing adversary-dilution checks. The following summarizes the main components.

Appendix A.1.1. Complexity of Building $A^{(h)}$

The h-hop adjacency matrix $A^{(h)}$ is computed via repeated sparse Boolean matrix multiplications:

- Define $B = A$.
- for $k = 1, \ldots, h$:
  - Update $A^{(h)} \leftarrow A^{(h)} \vee (B \neq 0)$.
  - Update $B \leftarrow \mathrm{sgn}(BA)$.

If the original adjacency matrix A has $M = \|A_0\|$ nonzeros, then computing all products up to hop h costs approximately:

$$O\left(\sum_{k=1}^{h} \left\| A^k \right\|_0 \right).$$

For sparse networks and moderate hop counts, this typically behaves as:

$$O(hM).$$

Thus, the multi-hop adjacency expansion is efficient and scalable for practical network sizes.

Appendix A.1.2. Degree and Trim-Budget Computation

Once $A^{(h)}$ is constructed:

- Computing the degrees $\deg_i^{(h)} = \sum_j A_{ij}^{(h)}$ costs

$$O\left(\left\| A^{(h)} \right\|_0 \right).$$

- Computing the trim budgets

$$f_i^{(h)} = \min\left\{ f_{\max}, \left\lfloor \tfrac{\deg_i^{(h)} - 1}{2} \right\rfloor \right\}$$

costs $O(N)$ and is negligible compared to other operations.

Appendix A.1.3. Spectral Connectivity: Estimating $\lambda_2(L^{(h)})$

Estimating the algebraic connectivity $\lambda_2(L^{(h)})$ of the h-hop Laplacian relies on a small number of Lanczos iterations. This typically requires:

$$O\left(\left\| A^{(h)} \right\|_0 \right).$$

 doi:10.20944/preprints202602.1237.v1

where t (often 10–50) is the number of iterations required for a stable eigenvalue estimate. This step generally dominates the proxy checks for moderate hop counts.

Appendix A.1.4. Adversary-Dilution Check

When adversary masks are known, the adversary-dilution condition:

$$\frac{\left|N_i^{(h)} \cap A\right|}{\left|N_i^{(h)}\right|} \leq \rho_{\max}$$

requires iterating over neighborhoods in $A^{(h)}$. The total cost is:

$$O\left(\left\|A^{(h)}\right\|_0\right).$$

This is lightweight relative to spectral computations.

Appendix A.1.5. Exact Robustness Checking

Exact checks for (r, s)-robustness are known to be computationally expensive, often requiring combinatorial or mixed-integer formulations. These are practical only for small to medium-sized graphs. This motivates the hybrid approach of combining exact checks (when feasible) with scalable proxies.

Appendix A.1.6. Summary of Complexity

For practical settings—including moderate hop counts and sparse graphs—the overall complexity of hop selection is dominated by:

- multi-hop adjacency construction: $O(hM)$,
- spectral estimator: $O(\|A^{(h)}\|_0 \cdot t)$.

These operations are significantly cheaper than full robustness checking and are therefore suitable for real-time or design-time resilient consensus applications.

References

1. LeBlanc, H.J.; Zhang, H.; Koutsoukos, X.; Sundaram, S. Resilient asymptotic consensus in robust networks. *IEEE Journal on Selected Areas in Communications* **2013**, *31*, 766–781.
2. Yuan, L.; Ishii, H. Resilient consensus with multi-hop communication. *IEEE Transactions on Automatic Control* **2025**.
3. Qin, J.; Ma, Q.; Shi, Y.; Wang, L. Recent advances in consensus of multi-agent systems: A brief survey. *IEEE Transactions on Industrial Electronics* **2016**, *64*, 4972–4983.
4. Amirkhani, A.; Barshooi, A.H. Consensus in multi-agent systems: a review. *Artificial Intelligence Review* **2022**, *55*, 3897–3935.
5. Olfati-Saber, R.; Fax, J.A.; Murray, R.M. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE* **2007**, *95*, 215–233.
6. LeBlanc, H.J.; Zhang, H.; Sundaram, S.; Koutsoukos, X. Resilient continuous-time consensus in fractional robust networks **2013**. pp. 1237–1242.
7. Saldana, D.; Prorok, A.; Sundaram, S.; Campos, M.F.; Kumar, V. Resilient consensus for time-varying networks of dynamic agents **2017**. pp. 252–258.
8. Dibaji, S.M.; Ishii, H. Resilient consensus of second-order agent networks: Asynchronous update rules with delays. *Automatica* **2017**, *81*, 123–132.
9. Ishii, H.; Wang, Y.; Feng, S. An overview on multi-agent consensus under adversarial attacks. *Annual Reviews in Control* **2022**, *53*, 252–272.
10. Zhao, D.; Lv, Y.; Yu, X.; Wen, G.; Chen, G. Resilient consensus of higher order multiagent networks: An attack isolation-based approach. *IEEE Transactions on Automatic Control* **2021**, *67*, 1001–1007.
11. Cao, Y.; Yu, W.; Ren, W.; Chen, G. An overview of recent progress in the study of distributed multi-agent coordination. *IEEE Transactions on Industrial informatics* **2012**, *9*, 427–438.

12. Pan, L.; Shao, H.; Mesbahi, M.; Xi, Y.; Li, D. Consensus on matrix-weighted switching networks. *IEEE Transactions on Automatic Control* **2021**, *66*, 5990–5996.
13. Usevitch, J.; Panagou, D. Resilient leader-follower consensus to arbitrary reference values in time-varying graphs. *IEEE Transactions on Automatic Control* **2019**, *65*, 1755–1762.
14. Shang, Y. Resilient consensus in multi-agent systems with state constraints. *Automatica* **2020**, *122*, 109288.
15. Fazeli, A.; Jadbabaie, A. Consensus over martingale graph processes. In Proceedings of the 2012 American Control Conference (ACC). IEEE, 2012, pp. 845–850.
16. Shang, Y. Median-based resilient consensus over time-varying random networks. *IEEE Transactions on Circuits and Systems II: Express Briefs* **2021**, *69*, 1203–1207.
17. Yemini, M.; Nedić, A.; Goldsmith, A.J.; Gil, S. Characterizing trust and resilience in distributed consensus for cyberphysical systems. *IEEE Transactions on Robotics* **2021**, *38*, 71–91.
18. Feng, Z.; Hu, G. Attack-resilient distributed convex optimization of cyber–physical systems against malicious cyber-attacks over random digraphs. *IEEE Internet of Things Journal* **2022**, *10*, 458–472.
19. Rezaee, H.; Parisini, T.; Polycarpou, M.M. Almost sure resilient consensus under stochastic interaction: links failure and noisy channels. *IEEE Transactions on Automatic Control* **2020**, *66*, 5727–5741.
20. Zhang, H.; Fata, E.; Sundaram, S. A notion of robustness in complex networks. *IEEE Transactions on Control of Network Systems* **2015**, *2*, 310–320.
21. Tyra, A.; Li, J.; Shang, Y.; Jiang, S.; Zhao, Y.; Xu, S. Robustness of non-interdependent and interdependent networks against dependent and adaptive attacks. *Physica A: Statistical Mechanics and its Applications* **2017**, *482*, 713–727.
22. Yuan, L.; Ishii, H. Reaching Resilient Leader-Follower Consensus in Time-Varying Networks via Multi-Hop Relays. *arXiv preprint arXiv:2411.09954* **2024**.
23. Su, L.; Vaidya, N.H. Reaching approximate Byzantine consensus with multi-hop communication. *Information and Computation* **2017**, *255*, 352–368.
24. Shang, Y. Resilient leaderless and leader–follower consensus over random networks through ℓ-hop communication. *European Journal of Control* **2024**, *79*, 101075. https://doi.org/https://doi.org/10.1016/j.ejcon.2024.101075.
25. Abbas, W.; Laszka, A.; Koutsoukos, X. Improving network connectivity and robustness using trusted nodes with application to resilient consensus. *IEEE Transactions on Control of Network Systems* **2017**, *5*, 2036–2048.
26. Coppersmith, D.; Fleischer, L.; Hendrickson, B.; Pinar, A. A divide-and-conquer algorithm for identifying strongly connected components **2003**.