

Observability- and Bandwidth-Aware Constraint Scheduling for Trust-Weighted Multi-Robot Map Fusion over Duty-Cycle-Limited LoRa Links

Magy Elbanhawy

magy.elbanhawy@eng.psu.edu.eg

New Mansoura University

Walaa Saber

Port Said University

Research Article

Keywords: Multi-robot SLAM, Constraint scheduling, LoRa, Duty-cycle regulation, Information gain, Observability, ROS 2

Posted Date: October 20th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11243516/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Observability- and Bandwidth-Aware Constraint Scheduling for Trust-Weighted Multi-Robot Map Fusion over Duty-Cycle-Limited LoRa Links

Abstract

Collaborative simultaneous localization and mapping (SLAM) let a team of robots build a single map, but beyond reliable Wi-Fi coverage the robots have to exchange inter-robot constraints over low-power links such as LoRa. Under a 1% duty-cycle limit the mapping front-end produces candidates far faster than the radio can carry them, so constraints wait minutes before transmission, and trust weighting at the fusion server cannot recover airtime that was already spent on stale or wrong measurements. Building on trust-weighted fusion, this paper proposes Bandwidth-Aware Constraint Scheduling (BACS) and its observability-aware variant, BACS+. The transmitter predicts the arrival age and server-side trust of each candidate, admits candidates through a trust gate, and fills the legal airtime budget by information per unit airtime, with BACS+ favoring weakly constrained robot pairs; the temporal decay is set from the measured scheduling deferral. The evaluation proceeded in stages, and every design choice was fixed on development seeds before testing on unused ones. In a controlled comparison on fresh seeds against the earlier trust-weighted framework at its published parameters, BACS+ lowered the mean map-alignment error from 0.281 m to 0.205 m (27.1%, 480 paired runs, two to five robots, four loss conditions), and most of this gain came from avoiding oldest-first transmission. Ranking candidates by predicted trust raised the error by 28–47%, because a transmitter cannot recognize outliers whereas server-side age decay suppresses them. Relative to random ordering the benefit was small (2.4% and 7.6% on two fresh seed blocks). On two ROS 2 robots with 868 MHz radios, the median deferral of transmitted constraints was 142–166 s against 0.17 s of channel delay, which supports calibrating temporal trust to the queueing timescale.

Keywords: Multi-robot SLAM, Constraint scheduling, LoRa, Duty-cycle regulation, Information gain, Observability, ROS 2.

1 Introduction

Multi-robot SLAM combines locally built trajectories and maps into one representation of a shared environment. Visual and LiDAR pipelines now give each robot an accurate local estimate [1, 2], and collaborative systems add inter-robot loop closures with centralized, distributed or hybrid graph

optimization so that the local estimates agree with each other [3–7]. Field campaigns such as the DARPA Subterranean Challenge showed, however, that the quality of a team map depends as much on what the robots manage to communicate as on how well each robot localizes itself [8, 9].

That dependence is sharpest in Internet of Robotic Things deployments that operate outside reliable Wi-Fi coverage, for instance in large facilities, agricultural fields or disaster sites. Long-range, low-power links such as LoRa fit these settings, yet their low data rate and the regulatory limits on airtime rule out any exhaustive exchange of mapping data [10–13]. The mapping front-end generates candidate constraints far faster than this, so only a small fraction can ever be sent. The order in which they leave the robot then becomes part of the estimator rather than a networking detail underneath it [14–25].

A second line acts after arrival: robust and trust-weighted back-ends limit the influence of inconsistent or delayed constraints. The trust-factor fusion framework of Saber et al. (EMRMF) [26], from which the present work started, multiplies a geometric consistency term by a temporal decay term at the fusion server. In that framework each robot periodically selects inter-robot observations for LoRa transmission, but the selection rule was left unspecified, and the evaluation assumed a 200 ms packet interval with delays of a few seconds. Under the 1% duty-cycle limit of the EU868 sub-band, the same link carries about six packets per minute and constraint age is governed by minutes of queueing, not by radio latency. The third and most common practice leaves the order to a first-in-first-out (FIFO) queue. FIFO sends the oldest candidate first, which under temporal decay is the constraint the server will trust least; scarce airtime is thus spent on measurements that the back-end is about to discount. None of these approaches decides which already-generated constraint should occupy the next legal transmission slot.

This paper introduces Bandwidth-Aware Constraint Scheduling (BACS) and its observability-aware extension, BACS+, which supply the missing transmitter-side selection and recalibrate temporal trust to the queueing timescale. The work was developed as a sequence of experiments, each prompted by the outcome of the previous one, and every later design choice was fixed on development seeds before it was tested on unused seeds. The main contributions are as follows:

- A transmitter-side scheduler that predicts the arrival age and the server-side trust of each candidate before airtime is spent, uses predicted trust only as an admission gate, and ranks admitted candidates by information per unit airtime; BACS+ adds a pairwise observability term whose surrogate is checked against exact incremental information.
- A temporal-decay rule, $\gamma = \ln 2 / T_{defer}$, tied to the measured scheduling deferral. Over 30 seeds and teams of two to five robots it stays within 10% of a coefficient tuned by parameter sweep and lowers map-alignment error by 18% relative to a coefficient derived from odometry drift.
- A staged evaluation with six alternative ordering rules, three trust-weighted rankings, a residual outlier gate and eight communication conditions. It shows that oldest-first transmission is the

dominant failure mode and that ranking by predicted trust is harmful, because the transmitter cannot recognize outliers while the server’s age decay keeps them harmless.

- Two pre-registered tests on previously unused seeds: a confirmation against random ordering and a controlled comparison with EMRMF at its published parameters, in which BACS+ lowered map-alignment error by 27.1% and a stepwise decomposition attributed the gain mainly to the scheduling rule.
- A hardware study on two ROS 2 robots with 868 MHz LoRa radios that measures the components of constraint age and confirms that scheduling deferral, rather than channel delay, dominates it.

The rest of the paper is organized as follows. Section 2 reviews related work. Section 3 formulates the scheduling problem and presents BACS and BACS+. Section 4 describes the experimental setup and the sequence of experiments, reports the simulation and hardware results, and closes with an analysis of what they show and where they stop. Section 5 concludes the paper and outlines future work.

2 Related Work

The literature relevant to BACS is reviewed here under four headings, following the order in which data move through a collaborative SLAM pipeline: where fusion takes place (system architecture), how much information is exchanged (communication-aware estimation), which measurements are worth sending (information-based selection), and how the back-end treats measurements after they arrive (robust weighting, trust, and observability). Separating the work in this way shows that each group addresses a different stage of the pipeline and makes clear where transmitter-side constraint scheduling fits. Table 1 summarizes the comparison.

Multi-robot SLAM architectures. Collaborative SLAM has moved from monolithic centralized pipelines toward distributed and hybrid systems that keep local autonomy while limiting communication [3]. DOOR-SLAM, Kimera-Multi, LAMP 2.0, and Swarm-SLAM demonstrate outlier-resilient distributed loop closure, dense metric-semantic reconstruction, and sparse decentralized collaboration [4–7], and certifiably correct and incremental distributed back-ends such as those of Tian et al. and McGann and Kaess make the optimizations itself communication-efficient [27, 28]. More recent systems, including SlideSLAM, Swarm-LIO2, Multi-SLAM with lightweight predictive frontier exploration, and CoMA-SLAM, further reduce the exchanged representation or distribute map updates to improve scalability [18–20, 22].

Communication-aware estimation and exploration. A second line of work treats bandwidth as an explicit resource. Resource-aware loop-closure detection selects which data to exchange so that estimation value is preserved under a communication budget [14], and distributed information gathering coordinates sensing under inter-robot communication constraints [15]. Lightweight and compressed multi-robot SLAM methods reduce descriptor, submap, or map traffic [16, 17, 23], and Bayer and Faigl showed that exploration can be coordinated on real robot teams by exchanging positions only [21]. Zhao et al. recently

found that even the timing phase of fixed-rate broadcasts changes team-level staleness in bandwidth-limited exploration [25]. On the radio side, LoRa surveys, mobility-aware LoRa protocols, and the ETSI spectrum-access rules describe the rate and duty-cycle limits that such systems face [11–13]. These studies support treating transmission timing and content as estimator-level decisions, but none of them predicts the trust that an individual constraint will receive after arrival.

Table 1 Comparison of representative approaches with BACS+.

Approach	Reduces traffic	Explicit airtime or bandwidth budget	Selects items before transmission	Models message age	Trust or outlier weighting	Information-based value	Pairwise observability
DOOR-SLAM [4]	✓	✗	✗	✗	✓	✗	✗
Kimera-Multi [5]	✓	✗	✗	✗	✓	✗	✗
Swarm-SLAM [7]	✓	✓	✓	✗	✓	✓	✗
Resource-aware loop closure [14]	✓	✓	✓	✗	✗	✓	✗
MI data selection [24]	✓	✗	✓	✗	✗	✓	✗
Desynchronised scheduling [25]	✓	✓	✗	✓	✗	✗	✗
Position-only exploration [21]	✓	✗	✗	✗	✗	✗	✗
Trust-factor fusion [26]	✗	✗	✗	✓	✓	✗	✗
BACS+ (this work)	✓	✓	✓	✓	✓	✓	✓

Information-based measurement selection. Information-theoretic selection ranks candidate measurements by their expected reduction of uncertainty. Resource-aware SLAM and bandwidth-efficient collaborative SLAM use graph structure or mutual information to keep informative data and discard redundant data [14, 24], and graph-connectivity indices have been linked formally to classical optimality criteria for active graph SLAM [29].

Robust weighting, trust, and observability. Robust and distributed pose-graph optimizations limits the damage caused by outliers once measurements are available [4, 5, 27], and trust-factor weighting additionally discounts constraints according to their age on arrival [26]. Observability analysis of multi-robot relative pose estimation shows that estimation quality also depends on sensing geometry and on the inter-robot constraints already present [30]. BACS+ uses this result as a scheduling prior: a candidate for a robot pair that is already well constrained receives lower priority than an equally informative candidate for a weakly constrained pair. The limitation of receiver-side methods for the present problem is one of timing, because they act after the airtime has been spent.

Taken together, recent work shows a clear shift toward communication-efficient collaborative SLAM, but it leaves open the transmitter-side decision that BACS addresses:

3 Proposed Method

BACS is a scheduling layer that sits on each robot between inter-robot candidate generation and the radio. Once per scheduling window it replaces the unknown value of a candidate with two quantities the robot can compute locally: a predicted server-side trust, which decides whether the candidate may be sent at all, and a predicted information gain, which orders the admitted candidates per unit of airtime. BACS+ adds a pairwise observability term to that gain. Section 3.1 states the system model and the scheduling problem inherited from trust-weighted fusion. Section 3.2 describes the scheduler in six steps: pooling new and deferred candidates, predicting their arrival age, predicting their trust, admitting or expiring them, scoring their information, and packing the airtime budget. Section 3.3 covers the calibration of temporal decay and an optional residual gate that was added during the evaluation.

3.1 System model and problem formulation

Multi-robot graph model. Consider a team of N robots exploring a shared planar environment. Robot i maintains a local pose sequence $\mathbf{x}_i = \{\mathbf{x}_{i,0}, \mathbf{x}_{i,1}, \dots\}$ with $\mathbf{x}_{i,k} \in SE(2)$. Local odometry and loop-closure measurements form intra-robot edges. When two robots observe overlapping parts of the environment, the front-end generates an inter-robot candidate constraint c between nodes (i, k) and (j, l) , with measured relative transform $\mathbf{z}_{ij}$ and information matrix Ω_{ij} . The fusion server keeps a global factor graph and solves the weighted nonlinear least-squares problem.

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \sum_{(i,j) \in \mathcal{E}} \theta_{ij} \mathbf{e}_{ij}(\mathbf{x})^T \Omega_{ij} \mathbf{e}_{ij}(\mathbf{x}), \quad (1)$$

where $\mathbf{e}_{ij}$ is the pose-graph residual and $\theta_{ij} \in [0,1]$ is a dynamic trust weight applied to inter-robot measurements [26]; local odometry edges carry unit weight.

Server-side trust. Following EMRMF [26], the server scores each arriving constraint by

$$\theta_{ij} = \max\left(0, 1 - \left(\frac{\|\mathbf{e}_{ij}\|}{\tau_e}\right)^p\right) \exp(-\gamma \Delta t_{ij}), \quad (2)$$

where τ_e is the admissible residual, p shapes the spatial roll-off, γ is the temporal decay coefficient and Δt_{ij} is the age of the constraint on arrival. A small floor keeps every edge numerically present. Fig. 1 plots both factors. With $p = 3$ the spatial term keeps consistent measurements near full weight and removes inconsistent ones quickly. The temporal term is where the two regimes part: with the coefficient $\gamma = 0.10 \text{ s}^{-1}$ used by EMRMF, a constraint that waited the 150 s typical of a duty-cycle-limited queue keeps less than one millionth of its weight, so the choice of γ cannot be separated from the scheduling policy that produces those ages.




Fig. 1 Components of the server-side trust of Eq. (2). (a) Spatial term for three shaping exponents. (b) Temporal term for the EMRMF coefficient ($\gamma = 0.10 \text{ s}^{-1}$), a drift-derived coefficient and the deferral-derived coefficient of Eq. (9); the grey histogram shows the measured deferral of constraints transmitted on the hardware testbed.

Regulatory airtime. The regulatory ceiling and the team-level allocation are kept distinct. For a device with duty-cycle limit δ and scheduling window W , the regulatory ceiling is δW . When N robots share one logical channel, robot i receives a fraction α_i of it:

$$B_i(W) = \alpha_i \delta W, \quad \sum_i \alpha_i \leq 1, \quad (3)$$

with equal sharing $\alpha_i = 1/N$ in all experiments. The airtime of a packet follows the standard LoRa model [11, 12]: with symbol time $T_{\text{sym}} = 2^{SF}/BW$, a payload of PL bytes occupies $T_{\text{air}}(PL) = (n_{\text{pre}} + 4.25 + n_{\text{pay}}) T_{\text{sym}}$, where n_{pay} is the payload symbol count. T_{air} is a step function of payload length, so candidates differ in radio cost as well as in expected value. At $SF = 7$, $BW = 125 \text{ kHz}$, coding rate 4/5 and a 52-byte payload a packet takes about 102.7 ms; with $\delta = 1\%$ and $W = 60 \text{ s}$ the per-device ceiling is 0.6 s per minute.

Scheduling problem. Let $\mathcal{C}$ be the candidates available to robot i in a window, each with unknown downstream value $v(c)$ and radio cost $T_{\text{air}}(c)$. The ideal decision chooses $x_c \in \{0,1\}$ to maximise $\sum_c x_c v(c)$ subject to $\sum_c x_c T_{\text{air}}(c) \leq B_i(W)$, a zero-one knapsack. The difficulty is not the knapsack but $v(c)$, which the transmitter cannot observe. BACS therefore splits the value into two parts: whether the constraint will still be trusted when it arrives, and how informative it will be if it is.

3.2 Transmitter-side scheduling with BACS and BACS+

Fig. 2 shows the architecture. Each robot runs local SLAM, a candidate generator and a BACS or BACS+ scheduler; local SLAM and the server-side optimizer of Eqs. (1) and (2) are unchanged. Selected constraints travel over the shared duty-cycle-limited LoRa channel to the fusion server, which periodically republishes the fused estimate. Each robot uses the fused poses to predict the residual the server will see,

and uses the count of constraints already delivered for each robot pair in the observability term of BACS+. Fig. 3 details the processing inside one window.




Fig. 2 System architecture. Each robot schedules its inter-robot candidates with BACS or BACS+ before transmission over the duty-cycle-limited LoRa channel; the fusion server performs trust-weighted pose-graph optimisation and returns fused-map feedback to the schedulers.




Fig. 3 Processing in one scheduling window. Steps 1–3 predict arrival age and server-side trust, and steps 4–6 admit, score and pack candidates into the airtime budget. Deferred candidates return to the queue with an incremented deferral count, and the server feeds back fused poses and delivered-constraint counts.

Arrival age. A candidate reaches the server after it has waited for later windows, for the packets ahead of it in the current window, and for its own transmission, as sketched in Fig. 4. Its age on arrival is predicted as

$$\widehat{\Delta t}_c = k_c W + \sum_{q \in Q_c} T_{air}(q) + T_{air}(c) + \widehat{T}_{retry}(c), \quad k_c = \frac{\sum_{q \in Q_c} T_{air}(q)}{\delta W}, \quad (4)$$

where $\mathcal{Q}_c$ is the set of packets queued ahead of c and k_c is the expected number of whole windows still to wait. With loss probability $\hat{p}_L$ and acknowledgement timeout T_{to} , the retry term is $\hat{T}_{retry} = \hat{p}_L(T_{air} + T_{to})/(1 - \hat{p}_L)$. Because queue position depends on rank and rank depends on predicted age, the implementation makes two passes: the first assumes every candidate sits mid-budget, the second recomputes Eq. (4) with the order from the first. The term $k_c W$ matters most; without it trust would be calibrated to a sub-second timescale while candidates routinely wait for minutes.




Fig. 4 Composition of constraint age under a duty cycle. A candidate generated early in a window may wait several windows for a transmission opportunity; the airtime and channel delay that follow are about three orders of magnitude shorter.

Predicted trust. The server’s residual needs current fused estimates of both endpoints. The transmitter substitutes the most recent fused poses $\hat{\mathbf{x}}^g$ it has received, forms the provisional residual $\tilde{\mathbf{e}}_{ij} = \mathbf{z}_{ij} - h(\hat{\mathbf{x}}_i^g, \hat{\mathbf{x}}_j^g)$, and predicts the trust that Eq. (2) will assign:

$$\hat{\theta}_{ij} = \max\left(0, 1 - \left(\frac{\|\tilde{\mathbf{e}}_{ij}\|}{\tau_e}\right)^p\right) \exp(-\gamma \widehat{\Delta t}_{ij}). \quad (5)$$

Raw local odometry is deliberately not used for this residual, because accumulated drift would then dominate the very quantity meant to judge the measurement.

BACS utility. The exact information a candidate would add, $I_c = 1/2 \log \det(I + \Omega_c J_c \Sigma_{ij,t} J_c^T)$ with Jacobian J_c and joint marginal covariance $\Sigma_{ij,t}$, is available only at the server. The transmitter instead scores each candidate by a local surrogate divided by its airtime:

$$u_c = \frac{\hat{I}_c}{T_{air}(c)}, \quad \hat{I}_c = w_v \bar{v}_c + w_d (1 - \bar{d}_c) + w_l \bar{l}_c, \quad (6)$$

where $\bar{v}_c$ is the coverage novelty of the candidate’s location in a 0.5 m occupancy bitmap, $\bar{d}_c$ the normalised degree of the local graph node and $\bar{l}_c$ the normalised loop span, with $w_v = 0.50$, $w_d = 0.20$ and $w_l = 0.30$. The three terms reward constraints that cover new ground, attach to weakly connected nodes, or close long loops over which drift has built up.

BACS+ observability. Novelty says nothing about whether the relative transform of a robot pair is already pinned down. BACS+ adds a pairwise deficit that decays with the number n_{ij} of constraints already delivered for the pair, as reported by the server:

$$O_{ij} = \exp\left(-\frac{n_{ij}}{n_{ref}}\right), \quad \hat{I}_c^+ = \hat{I}_c + w_o O_{ij}. \quad (7)$$

O_{ij} equals one for a pair with no delivered constraints and falls towards zero as the pair becomes well constrained (Fig. 5a), so scarce airtime is spread across pairs. The principal setting is $w_o = 0.30$, $n_{ref} = 6$. In a two-robot team there is only one pair, and the term cannot change the ranking.

Selection. Multiplying trust and information, $\hat{\theta}_c \hat{I}_c / T_{air}(c)$, is the obvious first choice, but it ranks poorly: exponential decay lets $\hat{\theta}$ span several orders of magnitude while $\hat{I}$ stays in a narrow band, so the order collapses to trust alone, and trust-first ordering favours constraints that already agree with the robot's estimate. BACS therefore uses predicted trust only for admission:

$$S^* = \operatorname{argmax}_S \sum_{c \in S} \frac{\hat{I}_c^{(+)}}{T_{air}(c)} \quad \text{s.t.} \quad \hat{\theta}_c \geq \theta_{gate} \quad \forall c \in S, \quad \sum_{c \in S} T_{air}(c) \leq B_i(W), \quad (8)$$

where $\hat{I}_c^{(+)}$ stands for $\hat{I}_c$ in BACS and $\hat{I}_c^+$ in BACS+, and $\theta_{gate} = 0.05$. If no candidate passes the gate in a window, the gate is relaxed for that window so that the airtime is not wasted. Section 4.4F later confirmed this design experimentally: folding predicted trust into the ranking, even through a well-conditioned concave transform, increased the error by 28–47%.

Eq. (8) is solved greedily, as listed in Algorithm 1: admitted candidates are sorted by information density and taken until the budget is exhausted, at a cost of $O(n \log n)$ per window and under 1 ms in simulation. Candidates that are not chosen stay queued with their deferral count incremented, so their predicted age and trust are refreshed in the next window. A candidate is dropped when its predicted trust falls below 1.5 times the numerical floor or after 12 deferred windows. Because trust prediction and information ranking are kept apart, a geometrically implausible constraint cannot enter the radio queue merely because it looks informative.

Algorithm 1 BACS/BACS+ selection for robot i in one scheduling window

Procedure
Input: queue $\mathcal{P}$ (deferred candidates), new candidates $\mathcal{N}$, fused poses $\hat{\mathbf{x}}^g$, pair counts n_{ij} , budget $B_i(W)$ (Eq. 3)
Output: ordered transmission set S ; updated queue $\mathcal{P}$
1: $\mathcal{P} \leftarrow \mathcal{P} \cup \mathcal{N}$
2: for pass = 1,2: for each $c \in \mathcal{P}$ compute $\widehat{\Delta t}_c$ (Eq. 4), $\hat{\theta}_c$ (Eq. 5), $\hat{I}_c$ (Eq. 6) and, for BACS+, $\hat{I}_c^+$ (Eq. 7)
3: remove from $\mathcal{P}$ every c with $\hat{\theta}_c < 1.5 \theta_{floor}$
4: $\mathcal{A} \leftarrow \{c \in \mathcal{P} : \hat{\theta}_c \geq \theta_{gate}\}$; if $\mathcal{A} = \emptyset$ then $\mathcal{A} \leftarrow \mathcal{P}$
5: sort $\mathcal{A}$ by $\hat{I}_c^{(+)}/T_{air}(c)$ in descending order; $S \leftarrow \emptyset$, $b \leftarrow 0$
6: for c in sorted $\mathcal{A}$: if $b + T_{air}(c) \leq B_i(W)$ then $S \leftarrow S \cup \{c\}$, $b \leftarrow b + T_{air}(c)$
7: transmit S ; for each $c \in \mathcal{P} \setminus S$: $k_c \leftarrow k_c + 1$, drop c if $k_c \geq 12$

3.3 Temporal calibration and optional residual gate

Deferral-calibrated decay. The coefficient γ in Eqs. (2) and (5) sets how fast a constraint loses weight. An early derivation tied it to odometry drift, $\gamma_{drift} = \ln 2 \sigma_d \bar{v} / \tau_e$, which gives about 0.0333 s^{-1} for $\sigma_d = 0.08$, $\bar{v} = 0.30 \text{ m/s}$ and $\tau_e = 0.5 \text{ m}$. That value proved about seven times too large (Section 4.4B), because under a duty cycle the ageing process that matters is waiting for a transmission opportunity, not motion during sub-second propagation. The decay is therefore anchored to the queue:

$$\gamma_{defer} = \frac{\ln 2}{T_{defer}}, \quad (9)$$

where T_{defer} is a measured central value of the deferral distribution, so a constraint loses half its temporal weight after one typical deferral. With $T_{defer} \approx 155 \text{ s}$ this gives $\gamma \approx 0.0045 \text{ s}^{-1}$. T_{defer} depends on the candidate rate, the channel share and the radio configuration, and has to be measured per deployment; the values measured for the conditions studied here are given in Section 4.1.

Residual gate. The trust gate rarely removes anything, because before the fused frames of a pair are aligned the predicted spatial term cannot tell a wrong data association from a correct one. After a pair has k delivered constraints, however, the raw predicted residual becomes informative. The residual-gated variant BACS+-RG therefore also rejects a candidate when $\|\tilde{\mathbf{e}}_c\| > \tau_r$ and $n_{ij} \geq k$, with no fallback (Fig. 5b). Both parameters were chosen on development seeds before any test, giving $\tau_r = 1 \text{ m}$ and $k = 3$. The gate changes admission only; the ranking of Eq. (8) is unchanged.




Fig. 5 (a) Observability deficit of Eq. (7) for the principal and the aggressive setting. (b) Admission region of the optional residual gate: candidates of pairs with fewer than $k = 3$ delivered constraints are always admitted, and afterwards only those with a predicted residual of at most $\tau_r = 1 \text{ m}$

4 Results

The evaluation was carried out as a sequence of experiments in which each result determined the next question. Section 4.1 describes the simulator, the communication conditions, the seed protocol and the hardware; Section 4.2 defines the metrics and statistical tests; Section 4.3 traces the sequence of experiments from EMRMF to the final comparison. The simulation results follow in Section 4.4 and the hardware results in Section 4.5, and Section 4.6 closes with an analysis of the findings and their limitations.

4.1 Experimental and environment setup

A. Simulation environment. The simulator released with this paper implements the whole pipeline: ground-truth world generation, distance-dependent odometry drift, inter-robot candidate generation, LoRa airtime and loss, trust prediction, information and observability scoring, scheduling, and a weighted SE(2) Gauss–Newton pose-graph back-end. Robots sweep overlapping partitions of a $14.0 \text{ m} \times 10.5 \text{ m}$ environment in a boustrophedon pattern at 0.30 m/s , adding a pose node every 2 s . A candidate constraint is generated whenever two robots are within 3.0 m of each other, with Gaussian measurement noise, and 12% of candidates are outliers caused by wrong data association. The server applies Eq. (2) in every experiment, so the FIFO policy is the trust-weighted fusion pipeline of EMRMF [26] without any transmitter-side selection. Table 2 lists the principal parameters. Policy comparisons use 480 s sessions, the decay study 720 s and the surrogate check 240 s .

Table 2 Principal simulation configuration

Parameter	Value
Environment	$14.0 \text{ m} \times 10.5 \text{ m}$
Robot speed $\bar{v}$ / pose interval	$0.30 \text{ m/s} / 2.0 \text{ s}$
Odometry drift σ_d	0.08 m per metre travelled
Observation radius / outlier rate	$3.0 \text{ m} / 12\%$
Admissible residual τ_e / shaping exponent p	$0.5 \text{ m} / 3$
LoRa carrier / SF / bandwidth / coding rate (C0)	$868 \text{ MHz} / \text{SF7} / 125 \text{ kHz} / 4/5$
Duty-cycle ceiling δ / scheduling window W (C0)	$1\% / 60 \text{ s}$
Trust gate θ_{gate}	0.05
Surrogate weights w_v, w_d, w_l	$0.50, 0.20, 0.30$
Principal BACS+ setting	$w_o = 0.30, n_{ref} = 6$
Residual gate (BACS+-RG)	$\tau_r = 1 \text{ m}, k = 3$
Temporal decay (C0)	$\gamma_{defer} = \ln 2 / 155 \text{ s} \approx 0.0045 \text{ s}^{-1}$
Team sizes	$N = 2-5$

B. Communication conditions and seed protocol. Two sets of conditions were used (Table 3). Conditions C0–C3 vary the radio regime: C1 halves the duty cycle, C2 lengthens each packet by moving to spreading factor 9, and C3 adds 20% independent packet loss. In C1 and C2 the window was lengthened to 120 s and 200 s , since with 60 s windows the per-robot budget of the largest team was shorter than one packet. For each of these conditions T_{defer} was measured in pilot runs on development seeds and γ was set by Eq. (9). Conditions L0–L3, used only in the comparison with EMRMF, keep the nominal radio and vary the loss

process from none to 30% burst loss with 0.2 s of added delay. Seeds were divided into five blocks with fixed roles, also listed in Table 3.

Table 3 Communication conditions (top)

Condition	Change from nominal	Window W	Measured T_{defer}	γ (s ⁻¹)	Delivered per run
C0 nominal	—	60 s	155 s	0.0045	32
C1 lower duty cycle	$\delta = 0.5\%$	120 s	77.8 s	0.0089	16
C2 longer packets	SF9	200 s	88.3 s	0.0078	15–16
C3 lossy link	20% independent loss	60 s	105.4 s	0.0066	31.8
L0 / L1 / L2	0% / 10% / 30% independent loss	60 s	155 s	0.0045	—
L3	30% burst loss + 0.2 s delay	60 s	155 s	0.0045	—

Table 3 seed blocks (bottom)

Seed block	Seeds	Role
Development	0–9	Surrogate fidelity; selection of the trust-weighted variant and the residual-gate parameters
Held-out I	10–39	Comparison with FIFO, alternative ordering rules, decay study, observability-weight ablation
Held-out II	40–69	Conditions C0–C3, trust-weighted ranking, residual-gate test
Confirmation	70–99	Pre-registered confirmation against random ordering
Bridge	100–129	Pre-registered comparison with EMRMF, conditions L0–L3

C. Physical testbed. The physical experiments used two AgileX LIMO differential-drive robots (LIMO-01 and LIMO-02), each fitted with an EAI T-mini Pro 2D LiDAR, an Orbbec DaBai RGB-D camera and an Intel NUC i7 running Ubuntu 22.04 and ROS 2 Humble Table 4. The robots and the fusion server each carried an RYLR998 LoRa transceiver set to 868 MHz, SF7, 125 kHz and coding rate 4/5. The scheduler node sat between candidate generation and the radio; local SLAM and the server optimiser were left unchanged. A Vicon system recorded both robots at 100 Hz. One session of about two hours was logged per policy (FIFO, BACS and BACS+) on three consecutive days, and each session was split into ten runs of 713–721 s. The bags contain the Vicon poses, wheel odometry (50 Hz), LiDAR scans (10 Hz), TF and the scheduler log, which records for every candidate its generation, selection, transmission and reception times, together with channel, RSSI, SNR and payload.

Table 4 Hardware and configuration of the recorded physical sessions

Component	Configuration
Robots	2 × AgileX LIMO differential-drive platforms
Sensors	EAI T-mini Pro 2D LiDAR; Orbbec DaBai RGB-D camera
Computing and middleware	Intel NUC i7 per robot; Ubuntu 22.04; ROS 2 Humble
Radio	RYLR998 LoRa, 868 MHz, SF7, 125 kHz, CR 4/5, 52-byte payload
Duty cycle / window	1% / $W = 60$ s
Ground truth	Vicon motion capture, 100 Hz
Policies	FIFO, BACS, BACS+
Runs	10 per policy (one ~2 h session per policy), 713–721 s each

Component	Configuration
Scheduler parameters	trust gate 0.5; $w_o = 0.30$; $n_{ref} = 6$; temporal half-life 60 s

4.2 Evaluation metrics and statistical analysis

Two error metrics are reported because they respond to different things. **Map-alignment RMSE**, the primary metric, measures how far the maps of different robots disagree at ground-truth co-location pairs in overlapping regions, which is precisely the error inter-robot constraints are meant to remove. **Per-step pose RMSE** measures each robot’s trajectory error against ground truth once the local gauge is fixed; most of it comes from the robot’s own odometry, so it need not move when only the scheduling of inter-robot constraints changes. Several mechanism measures complement these: delivered constraints, the mean server trust of delivered constraints (trust yield), outliers transmitted, and the share of the server’s total trust weight carried by outliers. On hardware the age of each transmitted constraint is split into scheduling deferral (selection minus generation time) and channel delay (reception minus transmission time).

All policies run on the same generated world for a given seed and team size, so their results form pairs. Paired differences are tested with the two-sided Wilcoxon signed-rank test, which makes no normality assumption, and effect size is given as Cliff’s δ , which lies between -1 and 1 (negative when the first policy has the lower error). Means carry t-based confidence intervals, and Holm correction is applied within each stated family. When the question is whether two policies perform alike, a non-significant difference is not enough, and an equivalence test with two one-sided tests on the log-ratio of paired errors and a $\pm 10\%$ margin is reported instead. Percentage changes are computed from means, $100(\bar{a} - \bar{b})/\bar{b}$ with b the reference.

4.3 Sequence of experiments

Fig. 6 summarize how the study developed from EMRMF to the final comparison. EMRMF established that server-side trust with temporal decay helps when delays are a few seconds. Measuring constraint age under a duty cycle showed that queueing, not radio latency, sets the age, which motivated both transmitter-side scheduling and a new calibration of γ . The comparison of BACS with FIFO then raised the question of which element of BACS produces the gain; the alternative ordering rules answered that avoiding oldest-first transmission accounts for most of it, which in turn led to testing whether trust should also shape the ranking. It should not, and the reason, fresh outliers, motivated the residual gate. With all parameters frozen, BACS+ was confirmed on unused seeds and finally compared, under a controlled protocol, with EMRMF itself.




Fig. 6 Sequence of the study, from the trust-weighted fusion of EMRMF to the controlled comparison with it. Each box names the experiment, the seeds used and the outcome that led to the next step

4.4 Simulation results

A. Composition of constraint age. Under the duty-cycle limit, propagation contributes almost nothing to constraint age (Fig. 7). A 52-byte packet occupies 0.103 s of airtime and the largest extra channel delay tested was 0.2 s, while the typical scheduling deferral is about 155 s, some three orders of magnitude longer. The temporal term of Eq. (2) therefore still matters, but it has to be calibrated against the queue, not against propagation.




Fig. 7 Timescales of constraint age in the simulated duty-cycle regime (log scale). Scheduling deferral exceeds packet airtime and the largest tested extra channel delay by about three orders of magnitude

B. Effect of temporal-decay calibration. Five decay rules were compared with the BACS scheduler on seeds 10–39 for two to five robots Fig. 8. The deferral-derived coefficient of Eq. (9) could not be distinguished from the value of 0.0030 s^{-1} found by a parameter sweep, for either metric and at any team size (Holm-corrected $p \geq 0.27$). The swept value was slightly better, by 2–5% in map-alignment error;

pooled over the 120 seed–team pairs this 3.5% difference reached significance ($p = 0.014$), and an equivalence test placed the deferral-derived rule within 10% of the swept one (90% CI of the error ratio 1.01–1.09, $p = 0.013$). The swept coefficient is an oracle obtained on the evaluation simulator itself, whereas Eq. (9) needs only a measured quantity. The drift-derived coefficient is roughly seven times larger and raised map-alignment error by 10–26% relative to Eq. (9), up to 0.354 m against 0.261 m for five robots; pooled over all pairs Eq. (9) was 18% lower ($p = 6.6 \times 10^{-8}$, lower in 92 of 120 pairs). An online deferral-adaptive rule behaved worst: its coefficient varied by orders of magnitude, it delivered fewer constraints for small teams, and its alignment error was 7–41% higher, because it learns only from delivered packets, which are systematically fresher than the backlog. Pose RMSE separated the rules only for two robots.




Fig. 8 Effect of the temporal-decay rule on (a) map-alignment RMSE and (b) per-step pose RMSE for two to five robots (mean and 95% confidence interval over 30 seeds). The deferral-derived coefficient stays close to the swept value, while the drift-based and deferral-adaptive rules raise map-alignment error, most of all in larger teams

C. Fidelity of the information surrogate. A scheduler needs to know how much a candidate would add at the moment of scheduling. The converged graph is the wrong reference for that, since it already contains the candidate and its neighbors, and regions a candidate helped to correct then look uninformative because the correction has been made. The base surrogate correlates positively with this reference at every team size, and the observability term improves the correlation in all 40 paired runs (Table 5). A diagnostic that used the converged graph instead produced correlations of about -0.4 to -0.6 , a reminder of how the wrong temporal reference reverses the conclusion.

Table 5 Spearman correlation between surrogate and exact incremental information (mean $\pm$ standard deviation over seeds 0–9 per team size)

N	Base ρ	BACS+ ρ	Median $\Delta\rho$	Seeds improved	One-sided p
2	0.45 ± 0.13	0.62 ± 0.08	+0.17	10/10	9.77×10^{-4}
3	0.38 ± 0.20	0.61 ± 0.13	+0.23	10/10	9.77×10^{-4}
4	0.42 ± 0.11	0.55 ± 0.08	+0.13	10/10	9.77×10^{-4}
5	0.49 ± 0.11	0.60 ± 0.08	+0.10	10/10	9.77×10^{-4}

D. Effect of scheduling policy relative to FIFO. With the deferral-derived coefficient in place, FIFO, BACS and BACS+ were compared on seeds 10–39 Table 6. Both BACS variants almost halved map-alignment error for two and three robots, with large effect sizes, and remained significantly better for four robots. For five robots BACS lost significance ($p = 0.067$) while BACS+ kept a 28.7% reduction ($p = 7.30 \times 10^{-4}$, $\delta = -0.53$). The gains are broad rather than driven by a few runs: BACS+ beat FIFO in 28, 30, 22 and 23 of 30 seeds for two to five robots (Fig. 9). Per-step pose RMSE showed no significant difference at any team size (all $p > 0.11$; Fig. 10a), and every policy delivered exactly the same number of constraints (32, 24, 32 and 40; Fig. 10b). The alignment gains therefore come from which constraints are sent, not from how many.

Table 6 Map-alignment RMSE relative to FIFO (mean over seeds 10–39); p from paired Wilcoxon tests

N	FIFO (m)	BACS (m)	BACS+ (m)	BACS vs FIFO	BACS+ vs FIFO
2	0.334	0.175	0.180	-47.6%; $p = 3.73 \times 10^{-9}$; $\delta = -0.82$	-46.0%; $p = 4.66 \times 10^{-8}$; $\delta = -0.79$
3	0.436	0.222	0.227	-49.1%; $p = 3.73 \times 10^{-9}$; $\delta = -0.92$	-48.1%; $p = 1.86 \times 10^{-9}$; $\delta = -0.90$
4	0.276	0.207	0.217	-25.2%; $p = 9.52 \times 10^{-4}$; $\delta = -0.42$	-21.3%; $p = 0.0113$; $\delta = -0.32$
5	0.257	0.211	0.183	-17.9%; $p = 0.0667$ (n.s.); $\delta = -0.46$	-28.7%; $p = 7.30 \times 10^{-4}$; $\delta = -0.53$




Fig. 9 Paired per-seed reduction in map-alignment RMSE relative to FIFO over seeds 10–39. Points are individual seeds, the dashed line marks parity with FIFO, and labels give the number of seeds in which each policy beats FIFO




Fig. 10 (a) Per-step pose RMSE with 95% confidence intervals; no comparison with FIFO is significant. (b) Delivered constraints per run, fixed by the runtime budget and identical for all policies

E. Effect of the ordering rule. To isolate the element of BACS responsible for the gain over FIFO, four more rules were run with exactly the same budget, window, expiry and packing: newest-first (LIFO), uniformly random order, ranking by predicted trust alone, and ranking by information density without the trust gate (Table 7). Every rule except trust-only beat FIFO by a wide margin for two and three robots, so the main source of error under a duty cycle is oldest-first transmission. Freshness alone is not the answer, though: pooled over team sizes (exploratory), LIFO was 17% worse than BACS+ ($p = 9.6 \times 10^{-4}$) and trust-only ranking 31% worse ($p = 9.3 \times 10^{-8}$). Random order, information-only ranking and BACS were not significantly different from BACS+ after Holm correction at any team size. The trust gate itself removed candidates only in larger teams; BACS and information-only ranking chose identical constraints in 30, 26, 13 and 1 of 30 seeds for two to five robots.

Table 7 Map-alignment RMSE (m) of seven rules under an identical runtime budget (mean over seeds 10–39)

Rule	Ranking key	$N = 2$	$N = 3$	$N = 4$	$N = 5$
FIFO	oldest first	0.334	0.436	0.276	0.257
LIFO	newest first	0.204	0.231	0.268	0.276
Trust-only	predicted trust $\hat{\theta}$	0.197	0.337	0.294	0.346
Random	uniform random order	0.175	0.225	0.236	0.217
Info-only	$\hat{I}/T_{air}$, no gate	0.175	0.226	0.211	0.198
BACS	$\hat{I}/T_{air}$, trust gate	0.175	0.222	0.207	0.211
BACS+	$\hat{I}^+/T_{air}$, trust gate	0.180	0.227	0.217	0.183

F. Effect of trust-weighted ranking under communication stress. Section 3.2 argued that predicted trust should gate candidates without ranking them. This was put to a direct test with three rankings that combine trust and information in a well-conditioned way, $\hat{I}_c^\theta = 1/2 \log(1 + \theta_c(e^{2\hat{I}_c^+} - 1))$, which is the exact gain of Section 3.2 with the information matrix scaled by trust as in Eq. (1). TW-now uses trust at the current age, TW-arrival uses trust at the arrival age of Eq. (4), and TW-arrival-sub adds diminishing returns within each window. The primary hypothesis, TW-now against random ordering, was fixed on development

seeds, and all rules were then run on seeds 40–69 in conditions C0–C3 Fig. 11. The hypothesis was rejected in every condition: TW-now had 28–47% higher alignment error than random ordering (Holm-corrected $p \leq 0.012$) and 35–51% higher error than BACS+ ($p < 10^{-4}$). Adding trust to the information ranking raised the error by 23–38% in every condition; predicting trust at arrival added another 1–4%, and diminishing returns changed nothing. FIFO, LIFO and trust-only ranking stayed among the worst rules throughout, and BACS+ had the lowest mean error in all four conditions, although its lead over random ordering on these seeds (3–14%) was not significant in any single condition.




Fig. 11 Map-alignment RMSE of representative rules under conditions C0–C3 (mean and 95% confidence interval over seeds 40–69 and $N = 2$ –5); the trust-weighted ranking is the worst non-FIFO rule in every condition

Table 8 shows why. The trust-weighted rankings did what they were built to do: the mean server trust of delivered constraints rose from about 0.42 to 0.54–0.56 and the median arrival age fell from about 107 s to 76–83 s. Yet the transmitter cannot tell outliers from inliers, since its predicted spatial term is close to one for both. Under BACS, BACS+ and random ordering an outlier arrives old and the server gives it a trust of about 0.013–0.023; under the trust-weighted rankings, LIFO and trust-only ranking it arrives fresh and receives about 0.17–0.21. The outliers’ share of the total trust weight in the fused graph grows roughly eightfold, and within each seed and team size that share tracks the alignment error (Spearman $\rho = 0.61$ in C0, 0.57–0.69 across conditions). The server’s age decay had in effect been filtering outliers, and a transmitter that prefers fresh constraints undoes that protection.

Table 8 Outlier mechanism in condition C0 (seeds 40–69, all team sizes)

Rule	Outliers transmitted	Server trust per outlier	Outlier share of trust weight	Map alignment (m)	Runs > 0.5 m
FIFO	337	0.014	0.003	0.306	12%
LIFO	280	0.190	0.017	0.284	8%
Trust-only	241	0.171	0.024	0.320	18%
Random	289	0.023	0.005	0.223	2%
BACS	331	0.013	0.003	0.219	2%
BACS+	322	0.014	0.003	0.204	3%
TW-now	238	0.193	0.024	0.294	18%

Rule	Outliers transmitted	Server trust per outlier	Outlier share of trust weight	Map alignment (m)	Runs > 0.5 m
TW-arrival	236	0.214	0.026	0.306	21%

G. Effect of the residual gate. Since predicted trust cannot expose outliers, the residual gate of Section 3.3 was tried as a direct filter. Its parameters were chosen in two stages on development seeds. Stricter thresholds ($\tau_r \leq 0.75$ m) caught more outliers but also rejected up to 57% of valid constraints and raised the error, and weighting the ranking by predicted trust on top of the gate increased the error by about a third; the chosen setting was therefore $\tau_r = 1$ m, $k = 3$ without trust weighting. The pre-registered test on seeds 40–69 required BACS+-RG to beat BACS+ significantly in at least three conditions and to lose significantly in none. It did not: the gate was 1.8% better in C0 (Holm-corrected $p = 2.1 \times 10^{-4}$), indistinguishable in C1 and C2, and in C3 its median error was lower (63 seed–team pairs better, 27 worse) while its mean was 2.0% higher ($p = 0.0033$). As a safety filter it worked as intended, rejecting 57% of outliers while losing 0.9% of valid constraints and cutting the outliers’ share of airtime from 9.0% to 5.8%. It is therefore kept as an option rather than as part of the final method.

H. Effect of the observability weight. A more aggressive setting ($w_o = 0.60$, $n_{ref} = 5$) was compared with the principal one on seeds 10–39. The two were statistically indistinguishable (pooled -1.5% , $p = 0.43$; per team size between -5.8% and $+2.4\%$, all Holm-corrected $p = 1.0$). For two robots they chose identical constraints in 23 of 30 seeds, as expected when only one robot pair exists.

I. Pre-registered confirmation on unused seeds. With all parameters frozen, and with the residual gate having failed its criterion, the pre-registration fixed BACS+ as the final method and tested a single hypothesis on seeds 70–99: that BACS+ has lower map-alignment error than random ordering, pooled over C0–C3 and all team sizes (480 pairs). Fig. 12 report the outcome. BACS+ was 2.4% lower on average (mean paired difference -0.006 m, 95% CI -0.020 to $+0.007$ m, $p = 0.091$, lower in 263 of 480 pairs), so the hypothesis was not confirmed. The pre-specified equivalence test did succeed: the geometric-mean error ratio was 0.974 (90% CI 0.942–1.007, $p = 3.4 \times 10^{-4}$), inside the $\pm 10\%$ margin. Against FIFO, BACS+ lowered the error by 33.2%, 15.7%, 6.8% and 26.3% in C0–C3, significantly in every condition except C2. It was also lower than BACS in all four conditions, significantly in C1 (-8.2% , Holm $p = 0.0049$), and 5.0% lower pooled ($p = 0.0085$, exploratory). Pose RMSE did not differ from random ordering (pooled -2.5% , Holm $p = 0.16$). An exploratory split by team size hints that random order is competitive for two and three robots whereas BACS+ is lower for four and five (in C0, 0.183 m against 0.224 m and 0.195 m against 0.217 m), which fits the observability term acting only when several pairs exist; this pattern was not pre-registered.




Fig. 12 Pre-registered confirmation on unused seeds 70–99. (a) Map-alignment RMSE per condition (mean and 95% confidence interval over $N = 2-5$). (b) Geometric-mean ratio of BACS+ to random-ordering error with its 90% confidence interval, which lies inside the $\pm 10\%$ equivalence margin

J. Controlled comparison with EMRMF. The last experiment measured how far BACS+ moves beyond the framework it grew out of. EMRMF was implemented inside the same simulator with its published trust parameters ($\gamma = 0.10 \text{ s}^{-1}$, $p = 3$, $\tau_e = 0.5 \text{ m}$, trust floor 0.01), no trust gate and no observability term. Because EMRMF does not specify how constraints are selected for transmission, FIFO order was assumed. Both arms, together with four intermediate arms, consumed identical precomputed worlds, candidate pools and channel realisations on fresh seeds 100–129, for two to five robots and conditions L0–L3; the protocol was registered and the code hashed before the run. Table 9 and Fig. 13 give the results. BACS+ lowered the mean map-alignment error from 0.281 m to 0.205 m, a 27.1% reduction (paired Wilcoxon $p = 1.2 \times 10^{-33}$, $\delta = -0.45$, lower in 371 of 480 pairs), and the reduction held in every condition: 26.8%, 26.7%, 24.5% and 30.0% in L0–L3 (Holm-corrected $p \leq 1.6 \times 10^{-8}$). Runs with an alignment error above 0.5 m fell from 8.3% to 1.0%, and pose RMSE dropped by 6.4% ($p = 1.7 \times 10^{-5}$).

The intermediate arms show where the gain comes from (Fig. 13b). Replacing EMRMF’s coefficient with the drift-derived one while keeping FIFO changed little (0.281 to 0.268 m, $p = 0.15$). Moving to the deferral-derived coefficient, still under FIFO, made things worse (0.305 m, $p = 8.4 \times 10^{-8}$): a slow decay stops discounting the stale constraints that FIFO keeps sending. Changing only the ordering, from FIFO to BACS+ at the same coefficient, then cut the error by 32.8% ($p = 3.4 \times 10^{-44}$). Eq. (9) is therefore beneficial only together with a scheduler that avoids oldest-first transmission, and the scheduling rule carries the improvement over EMRMF. In the pre-registered control comparison BACS+ was also 7.6% lower than random ordering ($p = 0.023$). On these seeds and loss conditions, however, BACS without the observability term was 4.5% lower still than BACS+ ($p = 0.023$, exploratory), so the end-to-end effect of the observability term is not consistent across seed blocks. An earlier version of this comparison, which

assumed the drift-derived coefficient for EMRMF and reused seeds 70–99, gave 0.248 m against 0.199 m; it is superseded by the protocol reported here.

Table 9 Controlled comparison with EMRMF on fresh seeds 100–129 (map-alignment RMSE in metres, mean over $N = 2-5$). All arms share worlds, candidate pools and channel realisations

Arm	Order	γ (s ⁻¹)	L0	L1	L2	L3	Pooled	Runs > 0.5 m
EMRMF (published)	FIFO	0.10	0.277	0.277	0.275	0.295	0.281	8.3%
EMRMF, drift γ	FIFO	0.0333	0.266	0.267	0.271	0.270	0.268	5.4%
FIFO, deferral γ	FIFO	0.0045	0.305	0.306	0.302	0.306	0.305	9.6%
Random	random	0.0045	0.219	0.219	0.215	0.235	0.222	2.5%
BACS	Eq. (8)	0.0045	0.194	0.194	0.199	0.197	0.196	0.0%
BACS+	Eq. (8)	0.0045	0.203	0.203	0.208	0.206	0.205	1.0%




Fig. 13 Controlled comparison with EMRMF on seeds 100–129. (a) Map-alignment RMSE of all six arms under conditions L0–L3 (mean and 95% confidence interval over $N = 2-5$). (b) Stepwise decomposition from EMRMF with its published coefficient, through FIFO with the drift-derived and the deferral-derived coefficients, to BACS+

4.5 Hardware results

A. Composition of constraint age on the physical link. Table 10 and Fig. 14 decompose the age of every transmitted constraint in the three physical sessions. The median channel delay was 0.17 s for all three policies (mean 0.170 s, standard deviation 0.003 s), which is consistent with one 0.103 s packet airtime plus serial and processing overhead. The median scheduling deferral of transmitted constraints was 153.5 s for FIFO, 141.9 s for BACS, and 166.1 s for BACS+, so deferral exceeded channel delay by factors of 902, 835, and 973. This confirms on real hardware the premise of Eq. (9): constraint age is set by transmission opportunity and backlog, not by radio latency. The measured deferral (142–166 s) is also close to the value $T_{defer} \approx 155$ s used in simulation. The per-run mean deferral did not differ significantly between FIFO and BACS+ (paired Wilcoxon $p = 1.0$, $\delta = -0.06$), which is expected because the scheduler changes which constraints are sent, not how many transmission opportunities exist.

Table 10 Physical two-robot results under the 1% duty-cycle condition (ten runs per policy). Timing values are medians over all transmitted constraints; interquartile ranges in brackets

Metric	FIFO	BACS	BACS+
Transmitted constraints (per session)	141	139	137
Scheduling deferral (s)	153.5 [112.3, 225.3]	141.9 [103.9, 188.0]	166.1 [117.1, 220.7]
Channel delay (s)	0.170	0.170	0.171
Constraint age on arrival (s)	153.7	142.0	166.2
Deferral / channel delay	902	835	973
Candidates logged per run, mean (range)	208.6 (201–215)	205.6 (200–217)	209.0 (201–219)
Transmitted per run, mean (range)	14.1 (12–15)	13.9 (12–16)	13.7 (12–16)
Mean RSSI (dBm) / SNR (dB)	−94.3 / 4.8	−95.7 / 5.0	−95.2 / 4.4
Airtime used, % of duty-cycle allowance	10.1	10.0	9.8




Fig. 14 Physical constraint-age decomposition. (a) Distribution of the scheduling deferral of all transmitted constraints for each policy. (b) Median scheduling deferral and median channel delay (log scale); deferral exceeds channel delay by about three orders of magnitude for all policies

B. Radio link and airtime. About 205–209 candidate constraints were logged per run, and 12–16 of them were transmitted. Every transmitted packet was received by the server, with mean RSSI between −94.3 and −95.7 dBm and mean SNR between 4.4 and 5.0 dB. All packets carried 52-byte payloads, so each occupied 0.103 s of airtime. The airtime used was about 10% of the combined duty-cycle allowance of the two transmitters, which shows that in these runs the selection rate was limited by the scheduler’s admission rule (trust gate 0.5 in the ROS 2 implementation) rather than by the regulatory ceiling itself.

C. Scope of the physical evidence. The physical sessions validate the timing mechanism behind Eq. (9) on real radios, real robots, and a real ROS 2 software stack. They do not provide a physical map-alignment comparison: the recorded bags contain Vicon poses, odometry, LiDAR scans, TF, and the scheduler log, but not the fused poses published by the server, so the map-alignment RMSE of Section 4.2 cannot be computed from them. The map-alignment results of this paper are therefore simulation results, and a physical map-alignment study is identified as future work in Section 5. The recorded sessions also used an earlier configuration of the ROS 2 scheduler (Table 4); the node in the repository has since been aligned with the formulation of Section 3 (trust gate 0.05, information-density ranking and a 155 s temporal half-

life), together with a rerun protocol that logs the fused poses and interleaves the policies. All logs, CSV files, OS image and generation scripts are available at [https:// github.com/-----/BACS-ROS2-SLAM](https://github.com/-----/BACS-ROS2-SLAM).

4.6 Results analysis

Oldest-first transmission is the main failure mode. The largest and most repeatable effect in the study is the cost of FIFO transmission. Under a duty cycle, the oldest candidate is also the one the server will discount most, so FIFO spends the scarce airtime where it helps least. Every ordering rule that avoids this recovered most of the loss. BACS+ lowered the error by 7–33% relative to FIFO on unused seeds, and by 27.1% relative to EMRMF at its published parameters, and the stepwise comparison in Section 4.4J attributes the improvement over EMRMF to the scheduling rule rather than to the decay coefficient.

Trust belongs at the fusion stage. Three separate experiments point the same way: trust-only ranking, the three trust-weighted rankings and trust weighting on top of the residual gate were all worse than simple alternatives. This supports the division of labor used in BACS, with trust applied at the fusion server as in EMRMF [26] and the transmitter restricted to admission and information-based ordering.

Information ranking and observability. The observability term consistently improved the fidelity of the information surrogate, but its end-to-end effect was mixed. BACS+ was lower than BACS on seeds 40–69 and 70–99, yet 4.5% higher on seeds 100–129 under loss conditions L0–L3, and relative to random ordering it was equivalent within 10% on one fresh block and 7.6% better on the other. Two features of the evaluated regime limit what any ranking can achieve: each robot sends only one or two constraints per window, which leaves little room to choose among non-stale candidates, and a two-robot team offers the pairwise term nothing to act on. The residual gate removed most harmful transmissions at almost no cost to valid ones, but at the simulated outlier rate the server’s trust weighting already kept outliers largely harmless, so the gate did not lower the error.

Temporal weighting should follow the queueing timescale. Across 30 seeds and four team sizes, the deferral-anchored coefficient stayed within 10% of a swept value, while drift-based coefficients raised the error by up to 30%, most for five robots. Simulation and hardware agree on the relevant timescale, about 155 s in simulation and 142–166 s on hardware against 0.1–0.2 s of radio delay. The comparison with EMRMF adds a qualification: a slow decay helps only when the scheduler stops sending stale constraints first, because under FIFO it removes the discounting that had masked FIFO’s weakness.

Scheduling affects map alignment far more than trajectory error. Inter-robot constraints mostly fix the relative placement of independently built local graphs, while the shape of each trajectory is dominated by local odometry. Pose RMSE showed no significant difference between BACS+, BACS, FIFO and random ordering, and only a 6.4% reduction relative to EMRMF, which makes map-alignment RMSE the appropriate end-to-end metric for transmitter-side scheduling.

5 Conclusion and Future Work

This work started from the trust-weighted fusion of EMRMF and a practical observation: when a duty-cycle-limited radio cannot carry every inter-robot constraint, weighting constraints at the server comes too late, because the airtime has already been spent. BACS moves the selection to the transmitter. It predicts the arrival age and server-side trust of each candidate, uses that trust as an admission gate, and allocates the legal airtime by information density, and BACS+ adds a preference for weakly constrained robot pairs. Four findings stand out. Oldest-first transmission is the dominant failure mode: BACS+ lowered map-alignment error by 7–33% relative to FIFO and, in a controlled comparison, from 0.281 m to 0.205 m (27.1%) relative to EMRMF at its published parameters, mainly through the scheduling rule. Predicted trust should gate but not rank transmissions, since trust-weighted rankings raised the error by 28–47% by letting fresh outliers through. Temporal trust should follow the queueing timescale; the deferral-derived coefficient stayed within 10% of a swept value, and on hardware the deferral of transmitted constraints (142–166 s) exceeded the channel delay (0.17 s) by about three orders of magnitude. Finally, the benefit of information-based ranking over random ordering in this regime is small, between parity and 7.6% on the two fresh seed blocks.

Several directions follow. The physical study should record the fused poses published by the server, run the aligned scheduler in a regime where the duty cycle is binding, and include three or more robots. The tendency of BACS+ to gain over random ordering in larger teams should be tested directly with more robots and less uniform overlap. The pair-count observability term could be replaced by a graph-derived measure of pairwise determinacy, such as a connectivity- or covariance-based summary broadcast by the server [28, 29]. An online estimator of queue-wide deferral is needed, since the adaptive rule tested here saw only delivered packets. Lastly, the residual gate could be extended towards the per-robot reputation layer suggested for trust-weighted fusion by Saber et al. [26] to cope with faulty or adversarial robots.

References

- [1] Campos, C., Elvira, R., Rodríguez, J.J.G., Montiel, J.M.M., Tardós, J.D.: ORB-SLAM3: An accurate open-source library for visual, visual–inertial, and multimap SLAM. *IEEE Transactions on Robotics* 37(6), 1874–1890 (2021). <https://doi.org/10.1109/TRO.2021.3130201>
- [2] Xu, W., Cai, Y., He, D., Lin, J., Zhang, F.: FAST-LIO2: Fast direct LiDAR-inertial odometry. *IEEE Transactions on Robotics* 38(4), 2053–2073 (2022). <https://doi.org/10.1109/TRO.2022.3141876>
- [3] Lajoie, P.-Y., Ramtoula, B., Wu, F., Beltrame, G.: Towards collaborative simultaneous localization and mapping: a survey of the current research landscape. *Field Robotics* 2(1), 971–1000 (2022). <https://doi.org/10.55417/fr.2022032>
- [4] Lajoie, P.-Y., Ramtoula, B., Chang, Y., Carlone, L., Beltrame, G.: DOOR-SLAM: Distributed, online, and outlier resilient SLAM for robotic teams. *IEEE Robotics and Automation Letters* 5(2), 1656–1663 (2020). <https://doi.org/10.1109/LRA.2020.2967681>
- [5] Tian, Y., Chang, Y., Herrera Arias, F., Nieto-Granda, C., How, J.P., Carlone, L.: Kimera-Multi: Robust, distributed, dense metric-semantic SLAM for multi-robot systems. *IEEE Transactions on Robotics* 38(4), 2022–2038 (2022). <https://doi.org/10.1109/TRO.2021.3137751>

- [6] Chang, Y., Ebadi, K., Denniston, C.E., et al.: LAMP 2.0: A robust multi-robot SLAM system for operation in challenging large-scale underground environments. *IEEE Robotics and Automation Letters* 7(4), 9175–9182 (2022). <https://doi.org/10.1109/LRA.2022.3191204>
- [7] Lajoie, P.-Y., Beltrame, G.: Swarm-SLAM: Sparse decentralized collaborative simultaneous localization and mapping framework for multi-robot systems. *IEEE Robotics and Automation Letters* 9(1), 475–482 (2024). <https://doi.org/10.1109/LRA.2023.3333742>
- [8] Ebadi, K., Bernreiter, L., Biggie, H., et al.: Present and future of SLAM in extreme environments: The DARPA SubT Challenge. *IEEE Transactions on Robotics* 40, 936–959 (2024). <https://doi.org/10.1109/TRO.2023.3323938>
- [9] Tian, Y., Chang, Y., Quang, L., Schang, A., Nieto-Granda, C., How, J.P., Carlone, L.: Resilient and distributed multi-robot visual SLAM: Datasets, experiments, and lessons learned. In: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 11027–11034 (2023). <https://doi.org/10.48550/arXiv.2304.04362>
- [10] Gielis, J., Shankar, A., Prorok, A.: A critical review of communications in multi-robot systems. *Current Robotics Reports* 3(4), 213–225 (2022). <https://doi.org/10.1007/s43154-022-00090-9>
- [11] Sun, Z., Yang, H., Liu, K., Yin, Z., Li, Z., Xu, W.: Recent advances in LoRa: A comprehensive survey. *ACM Transactions on Sensor Networks* 18(4), Article 67 (2022). <https://doi.org/10.1145/3543856>
- [12] ETSI: EN 300 220-2 V3.3.1: Short Range Devices (SRD) operating in the frequency range 25 MHz to 1 000 MHz; Part 2: Harmonised Standard for access to radio spectrum for non-specific radio equipment (2025). https://www.etsi.org/deliver/etsi_en/300200_300299/30022002/03.03.01_60/en_30022002v030301p.pdf
- [13] Lanka, V.S.V.G.S., Kothandaraman, D.: Mobility-aware LoRa protocol for reliable IoT communications. *IEEE Access* 14, 5870–5889 (2026). <https://doi.org/10.1109/ACCESS.2025.3649413>
- [14] Tian, Y., Khosoussi, K., How, J.P.: A resource-aware approach to collaborative loop-closure detection with provable performance guarantees. *The International Journal of Robotics Research* 40(10–11), 1212–1233 (2021). <https://doi.org/10.1177/02783649211021869>
- [15] Viseras, A., Xu, Z., Merino, L.: Distributed multi-robot information gathering under spatio-temporal inter-robot constraints. *Sensors* 20(2), 484 (2020). <https://doi.org/10.3390/s20020484>
- [16] Han, J., Ma, S.: Distributed multi-robot SLAM algorithm with lightweight communication and optimization. *Electronics* 13(20), 4129 (2024). <https://doi.org/10.3390/electronics13204129>
- [17] Zhang, L., Deng, J.: Deep compressed communication and application in multi-robot 2D-lidar SLAM: An intelligent Huffman algorithm. *Sensors* 24(10), 3154 (2024). <https://doi.org/10.3390/s24103154>
- [18] Liu, X., Lei, J., Prabhu, A., Tao, Y., Spasojevic, I., Chaudhari, P., Atanasov, N., Kumar, V.: SlideSLAM: Sparse, lightweight, decentralized metric-semantic SLAM for multi-robot navigation. *IEEE Transactions on Robotics* 41, 6529–6548 (2025). <https://doi.org/10.48550/arXiv.2406.17249>
- [19] Zhu, F., Ren, Y., Yin, L., et al.: Swarm-LIO2: Decentralized efficient LiDAR-inertial odometry for aerial swarm systems. *IEEE Transactions on Robotics* 41, 960–981 (2025). <https://doi.org/10.48550/arXiv.2409.17798>
- [20] Athukorala, A., Lau, B.P.L., Pongsirijinda, K., Yuen, C., Tan, U.-X.: Multi-robot collaborative SLAM (Multi-SLAM) with distributed lightweight predictive frontier exploration (LPFE). *IEEE Robotics and Automation Letters* 11(2), 2274–2281 (2026). <https://doi.org/10.1109/LRA.2025.3645990>
- [21] Bayer, J., Faigl, J.: Decentralized multi-robot exploration under low-bandwidth communications. *Autonomous Robots* 50, Article 7 (2026). <https://doi.org/10.1007/s10514-025-10234-3>
- [22] Chen, L., Su, Y., Wang, J., Han, P., Xia, Z., Bu, S., Li, K., Hu, B., Meng, S., Wang, G.: CoMA-SLAM: Collaborative multi-agent Gaussian SLAM with geometric consistency. *Proceedings of the AAAI Conference on Artificial Intelligence* 40(4), 2922–2929 (2026). <https://doi.org/10.1609/aaai.v40i4.37283>
- [23] Pu, F., Nawaz, M.W., Imran, M.A., Popoola, O.: ViT–NeRF edge-assisted multi-robot SLAM for scalable low-bandwidth exploration. In: *IEEE International Conference on Communications Workshops (ICC Workshops)* (2026). <https://doi.org/10.1109/ICCWorkshops63917.2026.11586277>

- [24] Zhang, W., Cheng, L., Wei, J., Xu, X., Zhang, Z., Xin, B.: Bandwidth-efficient collaborative SLAM: Mutual information for data selection and finite-state entropy for lossless compression. *IEEE Transactions on Automation Science and Engineering* 23, 5401–5412 (2026). <https://doi.org/10.1109/TASE.2026.3666939>
- [25] Zhao, R., Que, H., Sun, J., Xie, J., Zhang, Q., Qiao, F.: Don't fire together: Desynchronised communication scheduling for bandwidth-limited multi-robot exploration. *Sensors* 26(15), 4778 (2026). <https://doi.org/10.3390/s26154778>
- [26] Saber, W., El Banhawey, M., Khaled, A., Moawad, I., Rizk, R.: Enhancing multi-robot map fusion in ROS 2 via dynamic trust factor weighting. *Journal of King Saud University Computer and Information Sciences* 38, 644 (2026). <https://doi.org/10.1007/s44443-026-01079-6>
- [27] Tian, Y., Khosoussi, K., Rosen, D.M., How, J.P.: Distributed certifiably correct pose-graph optimization. *IEEE Transactions on Robotics* 37(6), 2137–2156 (2021). <https://doi.org/10.1109/TRO.2021.3072346>
- [28] McGann, D., iMESA: Incremental distributed optimization for collaborative simultaneous localization and mapping. In: *Robotics: Science and Systems (RSS)* (2024). <https://doi.org/10.48550/arXiv.2406.07371>
- [29] Placed, J.A. A general relationship between optimality criteria and connectivity indices for active graph-SLAM. *IEEE Robotics and Automation Letters* 8(2), 816–823 (2023). <https://doi.org/10.48550/arXiv.2110.01289>
- [30] Shin, K., Sim, H., Nam, S., Kim, Y., Hu, J., Kim, K.-K.K.: Multi-robot relative pose estimation in SE(2) with observability analysis: A comparison of extended Kalman filtering and robust pose graph optimization. *IEEE Transactions on Intelligent Vehicles* (2024). <https://doi.org/10.1109/TIV.2024.3478290>

Declarations

Funding. The authors received no specific funding for this work.

Competing interests. The authors declare no competing interests.

Ethics approval. Not applicable; no human participants or animal subjects were involved.