Mobility-Predictive Risk-Budgeted Protection Selection for Federated Learning in Space–Air–Ground Mobile Networks

Wenyu Ma
Information Engineering University
Zhengzhou, China
mwy_ndsc@163.com

Tao Hu
Information Engineering University
Zhengzhou, China
hutaondsc@163.com

Xinlei Liu
Information Engineering University
Zhengzhou, China
liuxl@stu.ouc.edu.cn

***Abstract*—Space–air–ground mobile networks make federated learning (FL) vulnerable to heterogeneous resources, handovers, satellite outages, and privacy overhead. We propose MP-RBPS, a mobility-predictive and risk-budgeted protection-selection framework for deadline-constrained FL. It predicts link risk, estimates on-time delivery reliability, jointly admits clients and selects standard protection policies, and updates residual exposure budgets. Experiments with non-IID data show that MP-RBPS reduces timeout by 46.2% over a reactive baseline and by 69.1% over Fixed-SecAgg, improves successful delivery by 16.0 percentage points under satellite outage, and reduces transmitted bytes by 54.2%.**

***Index Terms*—Federated learning; space–air–ground networks; mobile edge computing; privacy-preserving learning; mobility prediction; deadline-aware scheduling.**

I. Introduction

Space–air–ground integrated networks connect terrestrial devices, UAVs, satellites, edge aggregators, and cloud coordinators for wide-area intelligence. Their heterogeneous links and mobility create long delays, intermittent visibility, handovers, and resource constraints. FL avoids uploading raw data [1], but updates can still leak information [2]. Differential privacy, secure aggregation, and authenticated encryption reduce exposure at additional computation, communication, synchronization, and accuracy cost [3]–[5]. A fixed strong policy can miss deadlines, while a lightweight policy may be inadequate on risky cross-domain links.

Existing protected FL and communication-efficient FL generally optimize privacy or transmitted bytes independently [6]–[8]. They rarely jointly model mobility risk, deadline reliability, protection overhead, and residual exposure. MP-RBPS fills this gap by selecting clients and policies together. Its contributions are: (i) a deadline-aware delivery and residual-risk model; (ii) a lightweight edge algorithm using mobility prediction, feasible-policy filtering, and utility-based admission; and (iii) evaluation under handover, satellite outage, and non-IID data.

II. System Model and Problem Formulation

We consider ground, UAV, and satellite-connected clients, regional edge aggregators, and a cloud coordinator. Client i has state $x_i^t = (R_i, d_i, p_i, h_i, e_i)$ containing rate, delay, packet loss, handover indicator, and energy. For policy q,

$$T_i^t(q) = T_{\text{train}} + T_{\text{proc}}(q) + B_i^t(q)/R_i + d_i + h_i \Delta d_h. \quad (1)$$

The effective loss is $\tilde{p}_i = p_i + h_i \Delta p_h$. Satellite clients additionally have outage probability p_{out}. A lightweight EWMA predictor estimates mobility-link risk: $\hat{r}_i^t = \lambda \hat{r}_i^{t-1} + (1-\lambda) r_i^{t-1}$. The calibrated timely-delivery score is

$$P_i^t(q) = \Pr[T_i^t(q) \le \tau](1-\tilde{p}_i)(1-p_{\text{out}})(1-\hat{r}_i^t). \quad (2)$$

Candidate policies combine pruning, quantization, AEAD, secure aggregation (SecAgg), and optional differential privacy (DP): q_1: uncompressed AEAD; q_2: 30% pruning/16-bit AEAD; q_3: 50% pruning/8-bit AEAD; q_4: q_3+SecAgg; q_5: q_4+DP. Each policy has size $B_i(q)$, energy $E_i(q)$, accuracy penalty $A(q)$, and residual exposure cost $R_{\text{priv}}(q)$. A client contribution score G_i combines data size, update norm, diversity, and domain fairness.

TABLE I
Candidate protection policies.

Policy	Pruning	Quantization	AEAD	SecAgg	DP
q_1	0%	32-bit	Yes	No	No
q_2	30%	16-bit	Yes	No	No
q_3	50%	8-bit	Yes	No	No
q_4	50%	8-bit	Yes	Yes	No
q_5	70%	8-bit	Yes	Yes	Yes

Let $z_{iq} \in \{0,1\}$ indicate admission and policy assignment. MP-RBPS maximizes

$$\sum_{i,q} z_{iq}[G_i P_i(q) - \alpha B_i(q) - \beta E_i(q) - \gamma A(q) - \delta R_{\text{priv}}(q)] \quad (3)$$

subject to one policy per admitted client, $P_i(q) \ge \theta$, resource and deadline constraints, $\sum_i z_{iq} \le K$, and residual budget $R_{\text{priv}}(q) \le \Gamma_i^t$. SecAgg is allowed only when its admitted group has at least $m_{\min}$ clients. The risk score is a policy-comparison metric, not a formal DP guarantee.

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

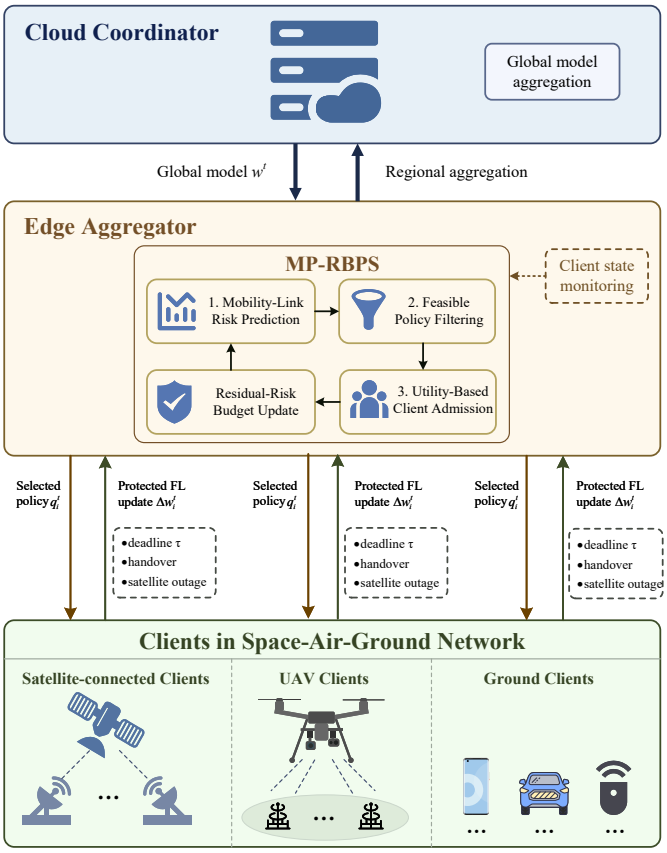


Fig. 1. Space–air–ground FL architecture and edge-side MP-RBPS control loop.

III. MP-RBPS Algorithm

At each round the edge aggregator (1) predicts $\hat{r}_i^t$ and computes G_i; (2) filters infeasible client–policy pairs using reliability, resource, energy, and budget constraints; (3) keeps the highest-utility feasible policy per client and admits the top-K clients; (4) checks the SecAgg group constraint and falls back to non-SecAgg policies if necessary; and (5) updates residual budgets after successful uploads. Complexity is $O(|\mathcal{D}||\mathcal{Q}| + |\mathcal{D}| \log |\mathcal{D}|)$, suitable for an edge aggregator.

The prediction is intentionally lightweight: the edge only keeps recent link outcomes, rather than raw training data. A high predicted-risk client is not automatically excluded. It remains eligible when compression or a different policy yields sufficient on-time reliability and utility. Conversely, a stable client may be rejected when its contribution is low or its residual budget is exhausted. This coupling prevents client selection from favoring only the best current links.

After a round, a successful protected update consumes the selected policy's exposure cost from Γ_i^t; failed or late updates are not aggregated. Repeated selection under weak protection therefore reduces future eligibility, whereas stronger policies preserve the residual budget at higher immediate cost. This is a system-level scheduling mechanism; DP-enabled policies still require a standard privacy accountant when formal privacy guarantees are reported.

TABLE II
Representative results.

Method	Acc.(%)	Timeout(%)	Success(%)	Comm.
FedAvg	62.1	–	–	1.00
Fixed-AEAD	63.4	53.0	42.0	1.03
Fixed-SecAgg	61.8	68.0	31.0	1.18
Compression-only	65.9	36.0	44.0	0.46
Reactive-policy	66.8	39.0	50.0	0.62
MP-RBPS	69.5	21.0	66.0	0.54

IV. Evaluation

We use 500 clients (50%/25%/25% ground/UAV/satellite), CIFAR-10 with Dirichlet non-IID partition ($\alpha = 0.5$), 25 clients per round, 150 rounds, three local epochs, and a 1.5-s deadline. Rates are Uniform$(8, 12)$, Uniform$(4, 8)$, and Uniform$(1, 4)$ Mbps for ground, UAV, and satellite clients; handover probability is swept to 0.5 and outage to 0.4. Each client trains a lightweight CNN using SGD (learning rate 0.01 and batch size 32). The wireless emulator samples domain-specific bandwidth, delay, packet loss, handover state, and satellite outage. All methods use the same emulator and three independent random seeds; reported curves and scalar values are means.

We compare MP-RBPS with FedAvg, Fixed-AEAD, Fixed-SecAgg, DP-FL, Compression-only, and Reactive-policy. Fixed approaches use one policy for all admitted clients. Reactive-policy observes the current link state but has neither mobility prediction nor a residual-risk budget. We report final test accuracy, timeout rate, successful update ratio, normalized communication and energy costs, and residual exposure cost. FedAvg and DP-FL are included as convergence and cost references when they are not part of a reliability sweep.

Table II summarizes representative results. MP-RBPS achieves the best accuracy and reliability trade-off. At handover probability 0.5, its timeout is 21.0%, versus 39.0% for Reactive-policy and 68.0% for Fixed-SecAgg. At outage probability 0.4, successful delivery is 66.0%, versus 50.0% and 42.0% for Reactive-policy and Fixed-AEAD. Fig. 2 shows that MP-RBPS reaches 69.5% accuracy, 2.7 points above Reactive-policy and 7.7 points above Fixed-SecAgg. The improvement results from preserving useful updates instead of applying costly protection uniformly. Fig. 3 confirms predictive selection under handover. Fig. 4 shows that the method retains valuable satellite clients when their expected contribution justifies the risk. Its normalized communication and energy costs are 0.54 and 0.71, compared with 1.18 and 1.25 for Fixed-SecAgg; ablations show that removing prediction raises timeout from 14.0% to 22.0%, while removing compression raises communication cost to 0.96.

A. Ablation and Domain Fairness

Fig. 5 quantifies the contribution of each component. Removing utility-based admission increases timeout to 26.0%; removing the residual-risk budget increases exposure cost from

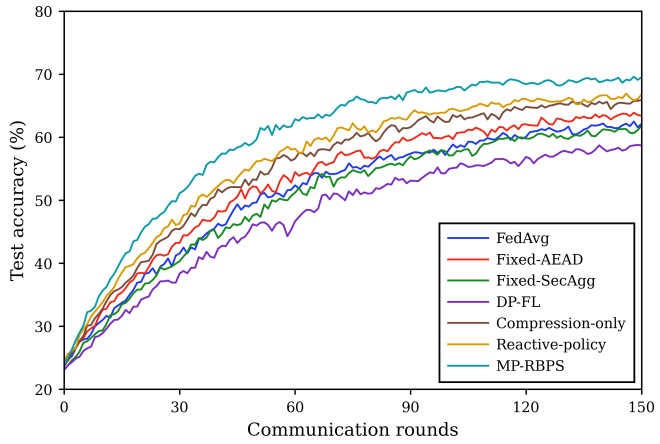


Fig. 2. Test accuracy versus communication rounds.

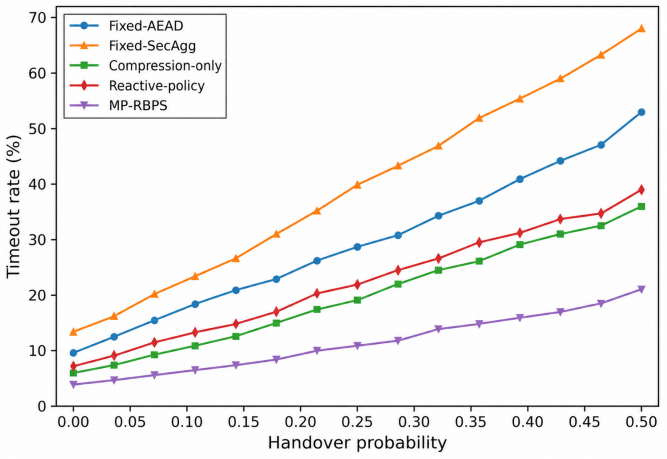


Fig. 3. Timeout rate versus handover probability.

0.42 to 0.61, and removing SecAgg raises it to 0.58. These results show that prediction improves reliability, compression preserves timeliness, and the budget controls repeated weak protection. The selected-policy distribution also reflects domain risk: ground clients mostly use lightweight policies, whereas UAV and satellite clients more frequently use SecAgg and DP-enabled policies. The fairness term preserves data-diversity coverage within 5.8% of the full-participation setting instead of simply over-selecting stable ground clients.

B. Operational Considerations

MP-RBPS uses only standard protection building blocks. AEAD provides confidentiality and integrity for each update in transit; the associated data should bind the client identifier, FL round, and sequence number, and nonce uniqueness must be enforced. AEAD alone does not prevent an authorized edge aggregator from seeing an individual decrypted update. SecAgg is therefore selected when aggregator-side exposure is significant and enough clients can form a valid group. DP is reserved for policies with high residual exposure, because clipping and noise add a measurable utility cost.

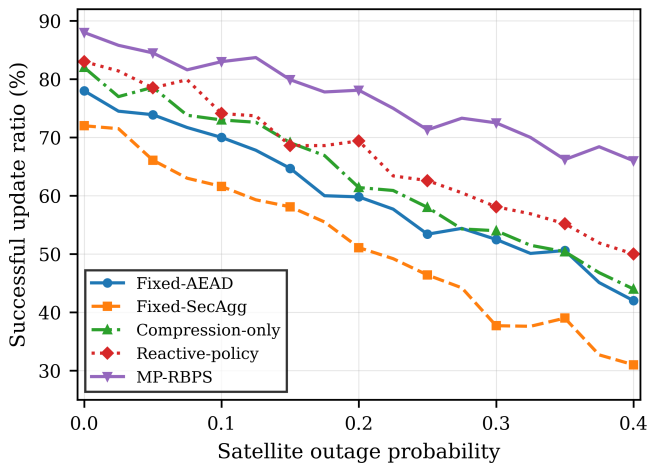


Fig. 4. Successful update ratio versus satellite outage.

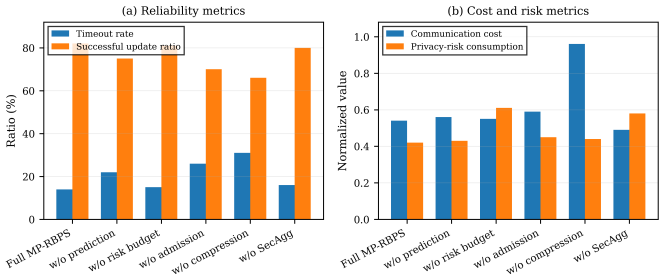


Fig. 5. Ablation results for MP-RBPS components.

The framework is deployable without retaining raw data at the edge. The predictor receives link measurements and past success outcomes; the contribution score can use local data size, prior update magnitude, privacy-preserving diversity summaries, and a domain-fairness weight. The candidate set is deliberately small, so the scheduler can re-evaluate all client–policy pairs at every round. In a production system, reliability scores should be calibrated from collected link traces, while formal DP accounting remains separate from the residual-risk score used for scheduling.

C. Limitations

The presented results use stochastic link emulation and a lightweight CNN; they establish the scheduling trade-off but do not replace trace-driven or device-level validation. The risk budget is an interpretable exposure-control heuristic rather than a cryptographic security parameter or an attack-success probability. Future experiments should incorporate LEO visibility traces, varying SecAgg dropout thresholds, heterogeneous device processors, and stronger gradient-inference attacks. These extensions will clarify how predictor calibration and policy choices transfer to operational space–air–ground deployments.

D. Discussion

The strongest available policy is not necessarily the best decision for every FL round. On a satellite or UAV link near a handover, extra cryptographic processing and synchronization may move an otherwise useful update beyond the deadline. Once discarded, that update contributes neither to learning nor to diversity. MP-RBPS instead compares the expected value of a protected update after accounting for its probability of timely delivery. It can choose a highly compressed policy for a time-critical client, defer a client whose predicted reliability is too low, or use SecAgg and DP when the expected contribution and available resources justify their overhead.

This decision rule also distinguishes communication protection from aggregator-side protection. A lightweight AEAD policy is appropriate when the main threat is wireless eavesdropping or tampering; a SecAgg policy is valuable when individual-update visibility at the aggregator is the relevant concern. The policy set thus represents practical implementation choices rather than an invented cryptographic primitive. Its individual settings, such as quantization width, pruning ratio, clipping bound, and SecAgg group size, can be adjusted to the capabilities and security requirements of a deployment.

Reliability prediction should be updated continuously. In the present simulation, the EWMA captures recent losses, timeout outcomes, handovers, and outages with low computational cost. An operational system may replace it with a calibrated logistic model or small neural predictor using visibility forecasts and radio measurements. Regardless of the predictor, its output should be treated as a calibrated scheduling score, not as an independent probability factor. Periodic recalibration against observed delivery outcomes is important when client mobility patterns or traffic loads change.

Finally, fairness is important in a space–air–ground setting. Selecting only low-risk terrestrial clients can improve immediate completion rate but degrade geographical coverage and data diversity. The contribution score therefore includes a domain-fairness term, allowing selected UAV and satellite clients to remain competitive when their information value offsets a moderate reliability penalty. The observed policy mix demonstrates this effect: stronger policies are used more frequently in higher-risk domains, while stable ground clients can use lighter policies without consuming the same resources.

E. Security Semantics and Reproducibility

The security objective of MP-RBPS is to reduce operational exposure while retaining useful updates under deadlines. It does not claim that compression itself provides privacy. Pruning and quantization only reduce payload size; their effects on information leakage require separate analysis. Likewise, the residual exposure cost ranks candidate policies under the stated threat model, but it must not be interpreted as a formal (ϵ, δ)-DP loss, a cryptographic bit-security level, or a probability that an inference attack succeeds. When DP is selected, clipping, noise multiplier, sampling rate, and composition must be recorded by a formal accountant independently of the scheduler.

For reproducible comparisons, all adaptive and baseline methods should share identical client partitions, model initialization, wireless traces, deadline, and client-admission budget. The policy overhead parameters should be reported explicitly, including protected-update size, encryption and SecAgg processing time, energy model, and any compression-induced accuracy penalty. Reliability results should distinguish the stress setting—for example, handover probability 0.5 or satellite outage probability 0.4—from the default training setting. This prevents a convergence result from being conflated with a worst-case delivery result and makes the trade-off in Table II interpretable.

In deployment, the edge aggregator can log only aggregate scheduling statistics: predicted-risk bins, selected policy counts, late-update counts, and group-formation failures. Such logs enable recalibration and audit without centralizing raw data or local model updates. A practical rollout can begin with AEAD and compression policies, then enable SecAgg only after stable client groups and dropout handling have been verified. DP-enabled policies should be activated according to the application threat model and an approved privacy budget, rather than solely because a link is unreliable.

V. Conclusion

MP-RBPS jointly adapts client admission and standard protection mechanisms to mobility, deadlines, resources, and residual exposure. It improves timely delivery and model accuracy while reducing communication and energy overhead in simulated space–air–ground FL. Future work will use trace-driven satellite mobility and real edge-device validation.

References

[1] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017, pp. 1273–1282.

[2] L. Zhu, Z. Liu, and S. Han, "Deep leakage from gradients," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019, pp. 14 774–14 784.

[3] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical secure aggregation for privacy-preserving machine learning," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2017, pp. 1175–1191.

[4] K. Wei, J. Li, C. Ma, M. Ding, W. Chen, J. Wu, M. Tao, and H. V. Poor, "Personalized federated learning with differential privacy and convergence guarantee," *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 4488–4503, 2023.

[5] National Institute of Standards and Technology, "Recommendation for block cipher modes of operation: Galois/counter mode (gcm) and gmac," NIST, Tech. Rep. SP 800-38D, 2007.

[6] M. Mansouri, M. Önen, W. B. Jaballah, and M. Conti, "Sok: Secure aggregation based on cryptographic schemes for federated learning," *Proceedings on Privacy Enhancing Technologies*, vol. 2023, no. 1, pp. 140–157, 2023.

[7] Y. Oh, N. Lee, Y.-S. Jeon, and H. V. Poor, "Communication-efficient federated learning via quantized compressed sensing," *IEEE Transactions on Wireless Communications*, vol. 22, no. 2, pp. 1087–1100, 2023.

[8] X. Chen, X. Zhou, H. Zhang, M. Sun, and H. V. Poor, "Client selection for wireless federated learning with data and latency heterogeneity," *IEEE Internet of Things Journal*, vol. 11, no. 19, pp. 32 183–32 196, 2024.