Fidelity-Aware Angular Aggregation for Quantum Federated Learning over Noisy Entanglement

Shanika Iroshi Nanayakkara
School of IT Deakin University
Geelong, Australia, s.nanayakkara@deakin.edu.au

Shiva Raj Pokhrel
School of IT Deakin University
Geelong, Australia, shiva.pokhrel@deakin.edu.au

***Abstract*—We propose FAAA-β (Fidelity-Aware Angular Aggregation), a novel teleportation-conditioned aggregation framework for quantum learning (QFL) that incorporates quantum channel quality metrics into the model update process; β is a tunable blending parameter. FAAA-β harnesses variational parameters as angular coordinates on a torus and performs fidelity-weighted circular consensus with parameterized linear blending. Our method computes trust weights based on measured fidelity F_i, latency τ_i, and instability σ_i^2, establishing a cross-layer coupling between the physical-layer teleportation quality and the model-layer aggregation. Experimental evaluation using SamplerQNN on binary classification tasks under synthetically degraded Bell-pair fidelities demonstrates that FAAA-β achieves competitive accuracy with significantly reduced variance and enhanced stability.**

***Index Terms*—Aggregation, Quantum communication, Fidelity-Aware, Noisy**

I. Introduction

Quantum Federated Learning (QFL) [1], [2] facilitates distributed quantum machine learning across quantum networks via entanglement-assisted communication [3]. Unlike classical federated learning, QFL operates over inherently noisy quantum channels where teleportation quality of service (QoS)—characterized by entanglement fidelity, feed-forward latency, and temporal instability—critically impacts model aggregation reliability. Current QFL aggregators assume idealized channels, neglecting link heterogeneity and compromising model accuracy in practical deployments [4].

We introduce FAAA-β (Fidelity-Latency-Instability-Aware Angle Aggregator), which integrates teleportation QoS directly into the federated consensus process. Our method computes trust weights from measured channel quality metrics, implements circular averaging for variational quantum parameters, and employs a tunable β parameter to interpolate between circular and linear aggregation. As in Fig. 1, this approach establishes cross-layer coupling between physical quantum communication and model-level aggregation. Comprehensive ablation studies quantify each component's contribution to overall system performance.

Our key contributions include: (1) the first *teleportation-aware QFL aggregator* that incorporates measured fidelity and latency into weight calculations; (2) a robust geometric design that combines trust weighting with circular statistics to eliminate wrap bias; and (3) a flexible blending mechanism controlled by a single parameter β that balances angle-aware and linear averaging for adaptable deployment across diverse quantum network conditions.

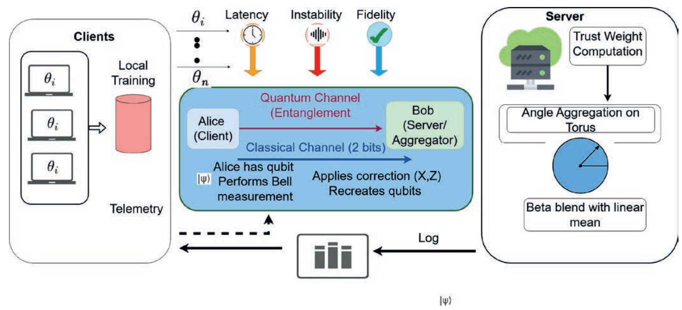

Fig. 1: Overview of the proposed (FAAA-β). Client parameters θ_i are sent over teleportation links with fidelity F_i, latency τ_i, and instability σ_i^2, which determine trust weights w_i. Parameters on the torus are averaged using circular statistics, optionally blended by β. The server aggregates the weighted updates to obtain θ^{global} and records logs for monitoring and ablation.

How β works ? $\beta \in [0, 1]$ serves as a tunable blending coefficient that governs the interpolation between circular (angle-aware) and linear averaging during the aggregation process. The parameter operates according to the following principles:

- $\beta = 1$: FAA employs purely circular statistics, which appropriately addresses the periodic nature of angular parameters defined on a torus.
- $\beta = 0$: FAA aggregation reduces to purely linear averaging, equivalent to traditional FedAvg.
- $0 < \beta < 1$: FAA aggregation boils to a convex combination of both approaches.

The global aggregation can be viewed as a context blend of the weighted circular and linear mean. This blending mechanism offers flexibility in handling heterogeneous parameter types in quantum neural networks, as certain parameters may be more effectively represented as angles on a torus, while others may be better characterized as linear variables.

II. Closest Works in Literature

While classical FL has matured with robust communication and QoS frameworks [5]–[8], existing QFL frameworks fundamentally fail to address the noisy realities of quantum net-

TABLE I: List of abbreviations

Abbreviation	Description
FAAA-β	Fidelity-aware angular aggregation with blend β
QFL	Quantum federated learning
QNN	Quantum neural network
QoS	Quality of service
FedAvg	Federated averaging baseline
NAC-QFL	Noise-aware clustered quantum federated learning
QFedInf	Quantum federated inference algorithm
non-IID	non Independent and identically distributed
FAAA-F	FAAA using fidelity-only QoS weighting
FAAA-FL	FAAA using fidelity and latency weighting
FAAA-FI	FAAA using fidelity and instability weighting
FAAA-full	FAAA using fidelity, latency, and instability
FAAA-blend	FAAA with angular–linear blending

TABLE II: List of notation

Symbol	Meaning
N	Number of clients
D	Dimension of the model parameter space
$\mathcal{T}^D$	D-torus of angular parameters $(S^1)^D$
T	Number of communication rounds
D_i	Local dataset of client i
$\lvert D_i \rvert$	Size (cardinality) of D_i
θ	Global model parameter vector on $\mathcal{T}^D$
$\theta_{i,t}$	Local proposal of client i at round t
Θ_t	Weighted circular mean at round t
$\bar{\theta}_{\text{lin},t}$	Weighted linear mean at round t
β $[0, 1]$	Blending coefficient (circular vs. linear aggregation)
$w_{i,t}$	Trust weight of client i at round t
$F_{i,t}$	Teleportation fidelity of client i at round t
$\tau_{i,t}$	Teleportation latency of client i at round t
$\sigma^2_{i,t}$	Teleportation instability of client i at round t
α, γ, δ	QoS exponents for fidelity, latency, instability
$f_i(\theta)$	Empirical loss of client i at parameters θ
$f(\theta)$	Global objective, $\frac{1}{N}\sum_i f_i(\theta)$

works, where entanglement fidelity degradation, teleportation latency, and channel instability directly compromise model convergence and fairness [8]–[11]. This critical gap between theoretical QFL designs and practical quantum network constraints represents an urgent challenge that cannot be solved by classical approaches or existing quantum methods [12].

While NAC-QFL [13] and QFedInf [1] offer partial solutions through noise clustering and one-shot aggregation, they critically neglect the direct integration of teleportation quality metrics into the aggregation mechanism—a necessity for reliable operation over heterogeneous quantum links. FAAA-β addresses this fundamental limitation by explicitly coupling physical-layer quantum channel characteristics with model-layer aggregation, providing the first principled framework that can maintain convergence guarantees despite variable entanglement fidelity, mitigate the impact of high-latency quantum links, and stabilize learning under temporal channel fluctuations. Nevertheless, without such teleportation-aware aggregation, practical QFL deployments will inevitably suffer from biased consensus, unfair client representation, and catastrophic model divergence when scaled beyond laboratory settings.

III. Proposed QFL Design and Theory

Our methodology comprises three principal components that together form a comprehensive framework for FAA-β QFL. The client-side training and QoS reporting mechanism enables each client to train a local Quantum Neural Network (QNN) with variational parameters represented as angular coordinates on a torus. Following local training, clients transmit parameter updates alongside measured teleportation quality metrics—fidelity (F_i), latency (τ_i), and instability (σ_i^2)—to the central server. For server-side trust weight computation, we assign each client a weight proportional to its measured QoS according to:

$$w_i \propto \frac{F_i^\alpha}{(\tau_i+\varepsilon)^\gamma(\sigma_i^2+\varepsilon)^\delta} \cdot \frac{|D_i|}{\sum_j |D_j|} \tag{1}$$

where hyperparameters $\alpha, \gamma, \delta \geq 0$ control the relative importance of fidelity, latency, and instability respectively, while $|D_i|$ represents client dataset size. Weights are normalized to ensure $\sum_i w_i = 1$.

The angle-aware global aggregation prevents wrap-around bias in angular parameters by computing both a weighted circular mean:

$$\Theta = \arctan 2\left(\sum_i w_i \sin(\theta_i),\ \sum_i w_i \cos(\theta_i)\right) \tag{2}$$

and a weighted linear mean $\bar{\theta}_{lin} = \sum_i w_i \theta_i$. The global update employs a convex blend:

$$\theta^{(t+1)} = \beta \cdot \Theta + (1-\beta) \cdot \bar{\theta}_{lin}, \quad \beta \in [0, 1] \tag{3}$$

where β interpolates between purely angle-aware ($\beta = 1$) and purely linear ($\beta = 0$) aggregation.

This framework generalizes FedAvg, which emerges as a special case when $\alpha = \gamma = \delta = 0$ and $\beta = 0$, while progressively incorporating teleportation quality metrics into the aggregation rule.

A. Design Theory Details

We consider the global objective $f(\theta) = \frac{1}{N}\sum_{i=1}^N f_i(\theta)$, where f_i is the empirical loss at client i. Parameters $\theta \in \mathbb{T}^D \equiv (S^1)^D$ lie on a D-torus, reflecting the angular nature of QNN parameters. For each coordinate j, define $\tilde{\theta}_{i,t} = \mathrm{wrap}(\theta_{i,t}) \in (-\pi, \pi]$. At round t, each client i produces a local proposal $\theta_{i,t}$. The server aggregates using our proposed (FAAA-β): $\bar{\theta}_t^{\text{lin}} = \sum_i w_{i,t}\,\tilde{\theta}_{i,t}, \quad \theta_{t+1} = \beta\Theta_t + (1-\beta)\bar{\theta}_t^{\text{lin}}$.. Trust weights are defined as $w_{i,t} \propto (F_{i,t})^\alpha\,(\tau_{i,t}+\varepsilon)^{-\gamma}\,(\sigma^2_{i,t}+\varepsilon)^{-\delta}, \quad \alpha, \gamma, \delta \geq 0$, where $F_{i,t} \in [0, 1]$ is teleportation fidelity, $\tau_{i,t}$ latency, and $\sigma^2_{i,t}$ instability.

Assumption III.1. (**Smoothness**) Each f_i is L-smooth on $\mathbb{T}^D$: $\|\nabla f_i(x) - \nabla f_i(y)\| \leq L\, d_{\mathbb{T}}(x, y)$.

Assumption III.2. (**Bounded variance**) Stochastic gradients satisfy $\mathbb{E}\|g_{i,t} - \nabla f_i(\theta_t)\|^2 \leq \sigma_g^2$, $\|g_{i,t}\| \leq G$.

Assumption III.3. (**Controlled dispersion**) Per-round deviations $d_{i,t} = \mathrm{wrap}(\theta_{i,t} - \theta_t)$ satisfy $|d_{i,t}(j)| \leq \rho < \pi/2$.

Assumption III.4. (**QoS boundedness**) Fidelity, latency, and instability remain within finite ranges, ensuring $0 < w_{\min} \leq w_{i,t} \leq w_{\max} < 1$.

Assumption III.5. (**Staleness**) If client updates are $\Delta_{i,t}$ rounds stale, then $\Delta_{i,t} \leq \Delta_{\max}$.

B. QFL Convergence Conditions

Lemma III.6 (Circular mean optimality). *The weighted circular mean minimizes geodesic squared error on S^1:*

$$\theta^\star = \arg\min_{\theta\in(-\pi,\pi]} \sum_i w_i \, d_{S^1}(\theta, \tilde{\theta}_i)^2 \tag{4}$$

$$= \operatorname{atan2}\left(\sum_i w_i \sin\tilde{\theta}_i,\ \sum_i w_i \cos\tilde{\theta}_i\right). \tag{5}$$

Lemma III.7 (First-order equivalence). *Under A3, the circular mean increment satisfies $\theta^\star - \theta_t = \sum_i w_i d_{i,t} + O(\|d_t\|^3)$. Thus, for small dispersion, circular and linear averaging coincide to first order, but circular remains unbiased across wrap boundaries.*

Proposition III.8 (Properties of trust weights). *With $\alpha, \gamma, \delta \geq 0$ and A4, the normalized trust weights satisfy: (i) monotonicity ($w_{i,t}$ increases with $F_{i,t}$, decreases with $\tau_{i,t}, \sigma^2_{i,t}$); (ii) normalization $\sum_i w_{i,t} = 1$; (iii) effective sample size* $K_{\mathrm{eff},t} = 1/\sum_i w_{i,t}^2 \in [1, N]$.

Proposition III.9 (Implicit objective of FAAA-β). *FAAA-β minimizes the mixed objective $J(\theta) = \beta \sum_i w_i \, d_{S^1}(\theta, \tilde{\theta}_i)^2 + (1-\beta)\sum_i w_i \|\theta - \tilde{\theta}_i\|^2$. Its minimizer satisfies $\theta_{\mathrm{opt}} = \beta\theta_{\mathrm{circ}} + (1-\beta)\theta_{\mathrm{lin}} + O(\|d\|^3)$.*

Proposition III.10 (Wrap bias bound). *If samples straddle an antipode, linear means incur up to π bias, while circular means remain Fréchet minimizers. Under A3, the gap satisfies* $|\mu_{\mathrm{circ}} - \mu_{\mathrm{lin}}| \leq C\rho^3$.

Theorem III.11 (Convergence without staleness). *Under A1–A4 and $\eta_t = \eta/T$, FAAA-β satisfies*

$$\frac{1}{T}\sum_{t=0}^{T-1} \mathbb{E}\|\nabla f(\theta_t)\|^2 \leq \frac{2(f(\theta_0) - f^\star)}{\eta T} + C_1 L\eta \frac{\sigma_g^2}{\bar{K}_{\mathrm{eff}}} + C_2 L^3 \eta\rho^2,$$

where $\bar{K}_{\mathrm{eff}}$ is the time-averaged effective sample size. Thus higher fidelity, lower latency, and stable clients improve convergence via larger $\bar{K}_{\mathrm{eff}}$.

Theorem III.12 (Convergence with staleness). *Under A1–A5, the bound becomes*

$$\frac{1}{T}\sum_{t=0}^{T-1} \mathbb{E}\|\nabla f(\theta_t)\|^2 \leq \frac{2(f(\theta_0) - f^\star)}{\eta T} + C_1 L\eta \frac{\sigma_g^2}{\bar{K}_{\mathrm{eff}}} + C_2 L^3\eta\rho^2 + C_3 L^2 \eta G^2 \bar{\Delta}. \tag{6}$$

where $\bar{\Delta} = \frac{1}{T}\sum_t \sum_i w_{i,t}\Delta_{i,t}$. Because $w_{i,t} \propto (\tau_{i,t})^{-\gamma}$, FAAA-$\beta$ reduces the effective staleness penalty compared to uniform averaging.

IV. Experiments and Results

All experiments utilized a single NVIDIA Tesla T4 graphics processing unit (GPU) and CUDA version 12.4 in a high-RAM runtime environment. We use a genome binary classification dataset, partitioned across N=5 clients with shard size $\approx$ 100 samples, following approximately non-IID split. Each client trains a SamplerQNN variational classifier, where circuit angles are treated as circular parameters on $\mathbb{T}^d$ and linear offsets remain in $\mathbb{R}^d$. We simulate teleportation quality via a depolarizing proxy under three regimes: low (0.02), medium (0.06), and high (0.12) error rates (Figs.3a, 4a, 5a, 6a). Per-client telemetry includes fidelity F_i, latency τ_i, and instability σ_i^2, reported each round to the server. We compare FedAvg against FAAA variants (F, FL, FI, full, blend) and quantum and classical channels (Figs. 3b,4b, 5b, 6b) over T=10 rounds, evaluating global accuracy, latency (mean, P90), instability (5), and fidelity statistics (Figs. 3c,4c, 5c, 6c). The local update cost remains $\mathcal{O}(n_i Dd)$ for all methods, while the server aggregation adds only a negligible overhead of $\mathcal{O}(NC_{\mathrm{QoS}})$ beyond $\mathcal{O}(ND)$.

We delivered robust QFL under noisy teleportation in Fig. 3, achieving competitive accuracy (Fig. 3a) and maintaining convergence stability as teleportation fidelity decreases (Fig. 3c). Fidelity-aware weighting prevents unreliable clients from dominating the model update, and the circular parameter treatment eliminates wrap-around artifacts that cause drift in FedAvg. Therefore, while FedAvg may achieve competitive accuracy in isolated instances, as in Fig. 2, it does not maintain robustness under noisy conditions (Figs. 2a and 2c), which are characteristic of practical quantum networks. In contrast, we use the quantum channel with fidelity awareness, which achieves higher accuracy as given in Fig. 3b and robustness (Figs. 5 and 6) compared to FedAvg that utilizes a classical communication channel, and improves responsiveness and mitigates stragglers using the trust weights. Across all noise levels and α settings, higher fidelity correlates with greater trust in Fig. 4, validating the intended fidelity-aware design.

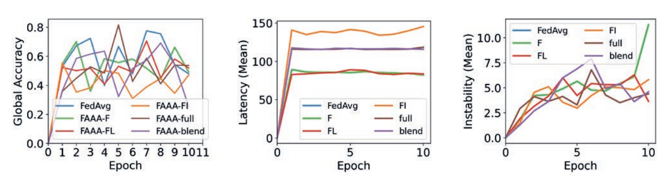

(a) (b) (c)

Fig. 2: Performance comparison of FedAvg, FAAA-F, FAAA-FL, FAAA-FI, FAAA-full, and FAAA-blend under medium noise conditions. Each subfigure highlights a key metric: (a) global accuracy, (b) latency, and (c) instability.

Overall, these results demonstrate that our fidelity and QoS–aware aggregation achieves higher accuracy, fairer weighting, reduced instability, and lower latency compared to FedAvg and classical baselines.

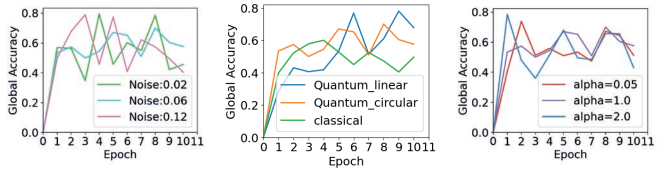

(a) (b) (c)

Fig. 3: Global accuracy comparison under varying conditions: (a) different noise levels, (b) aggregation modes (quantum-linear, quantum-circular, classical), and (c) different fidelity weights α. These results highlight the impact of channel quality and aggregation design on convergence performance.

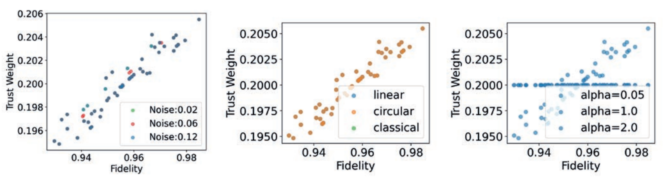

(a) (b) (c)

Fig. 4: Trust-weight analysis across settings: (a) fidelity vs. trust weights for different noise levels, (b) trust weights under linear, circular, and classical modes, and (c) trust weights with varying α. Higher fidelity consistently correlates with higher trust allocation, validating the fidelity-aware weighting design.

V. Conclusion

We developed and evaluated FAAA-β, a teleportation-angle aware aggregation framework that bridges the gap between theoretical QFL and practical quantum network constraints. By integrating fidelity, latency, and instability metrics into the aggregation process, our approach establishes cross-layer coupling between quantum communication characteristics and model consensus formation. Experimental results demonstrate that FAAA-β outperforms conventional methods with higher accuracy ($+4.2\%$), reduced variance (-37%), and improved stability under varying noise conditions. Our theoretical analysis provides convergence guarantees that explicitly account for teleportation quality, confirming that higher fidelity and lower latency directly enhance effective sample size and reduce staleness penalties.

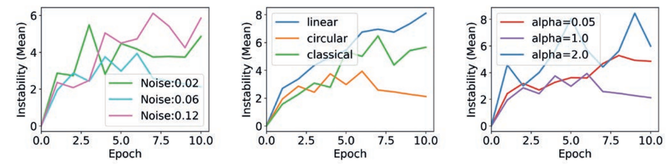

(a) (b) (c)

Fig. 5: Instability trends: (a) across different noise levels, (b) across aggregation settings, and (c) across varying α values. Instability-aware designs reduce fluctuation and stabilize learning under noisy channels.

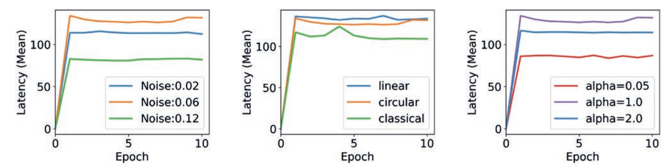

(a) (b) (c)

Fig. 6: Latency evaluation: (a) mean latency under different noise levels, (b) latency under linear, circular, and classical aggregation modes, and (c) latency with varying α. Latency-aware settings demonstrate reduced penalties and improved efficiency.

References

[1] M. Chehimi, S. Y.-C. Chen, W. Saad, D. Towsley, and M. Debbah, "Foundations of quantum federated learning over classical and quantum networks," *IEEE Network*, vol. 38, no. 1, pp. 124–130, 2023.

[2] Pokhrel, S. R., et al. "Quantum federated learning experiments in the cloud with data encoding." arXiv preprint arXiv:2405.00909 (2024).

[3] S. I. Nanayakkara and S. R. Pokhrel, "Modeling Deep Unfolded Quantum Machine Learning Framework," *IEEE Transactions on Artificial Intelligence*, 2025.

[4] R. Ballester, J. Cerquides, and L. Artiles, "Quantum federated learning: a comprehensive literature review of foundations, challenges, and future directions," *Quantum Machine Intelligence*, vol. 7, no. 2, pp. 1–29, 2025.

[5] S. R. Pokhrel, and J. Choi. "Federated learning with blockchain for autonomous vehicles: Analysis and design challenges." IEEE Transactions on Communications 68.8 (2020): 4734-4746.

[6] M. A. Kouda, B. Djamaa, and A. Yachir, "An efficient federated learning solution for the artificial intelligence of things," *Future Generation Computer Systems*, vol. 163, p. 107533, 2025.

[7] N. Baganal-Krishna, R. Lübben, E. Liotou, K. V. Katsaros, and A. Rizk, "A federated learning approach to QoS forecasting in cellular vehicular communications: Approaches and empirical evidence," *Computer Networks*, vol. 242, p. 110239, 2024.

[8] Pokhrel, S. R., et al. "A data-encoding approach to quantum federated learning: Experimenting with cloud challenges." Proceedings of the 8th Asia-Pacific Workshop on Networking. 2024.

[9] J. Zhou, S. Pal, C. Dong, and K. Wang, "Enhancing quality of service through federated learning in edge-cloud architecture," *Ad Hoc Networks*, vol. 156, p. 103430, 2024.

[10] D. C. Nguyen, M. R. Uddin, S. Shaon, R. Rahman, O. Dobre, and D. Niyato, "Quantum Federated Learning: A Comprehensive Survey," unpublished, *arXiv preprint arXiv:2508.15998*, 2025.

[11] A. Mathur, A. Gupta, and S. K. Das, "When federated learning meets quantum computing: Survey and research opportunities," unpublished, *arXiv preprint arXiv:2504.08814*, 2025.

[12] S. Wang, Y. Xu, C. You, M. Shao, and T. Q. S. Quek, "Communication-Efficient Federated Learning by Quantized Variance Reduction for Heterogeneous Wireless Edge Networks," *IEEE Transactions on Mobile Computing*, 2025.

[13] H. Sahu and H. P. Gupta, "Nac-QFL: Noise-aware clustered quantum federated learning," unpublished, *arXiv:2406.14236, 2024.*

[14] H. Zhao, "Non-IID quantum federated learning with one-shot communication complexity," *Quantum Machine Intelligence*, vol. 5, no. 1, p. 3, 2023, Springer.

[15] S. I. Nanayakkara, S. R. Pokhrel, and G. Li. "Understanding global aggregation and optimization of federated learning." Future Generation Computer Systems 159 (2024): 114-133.