Utility Aware Adaptive Privacy Budget Allocation for Streaming Multi-Agent Systems

Puspanjali Ghoshal
Indian Institute of Technology
Guwahati, India
g.puspanjali@iitg.ac.in

Ashok Singh Sairam
Indian Institute of Technology
Guwahati, India
ashok@iitg.ac.in

ABSTRACT

Streaming Multi-Agent Systems (MAS) involve autonomous agents continuously sharing observations with a fusion center for coordination. In such scenarios, the data flow is mostly consistent allowing stronger privacy preservation at the cost of reduced utility. However, anomalies like accidents or emergencies require timely and accurate information sharing, necessitating lower privacy. Applying conventional differential privacy with a fixed and low ϵ value (high noise), risks obscuring the critical information. As each report consumes a part of the agents' differential privacy budget, by sequential composition, the overall privacy reduces over time. Premature depletion of the privacy budget will lead to lower available budget for periods of high variability, reducing the accuracy and effectiveness of long term analytics. In this paper, we propose Adaptive Privacy Budget Allocation (APBA), a dynamic privacy budget allocation mechanism, leading to a tradeoff between utility and privacy. The allocation depends on (i) local signal uncertainty and (ii) the agent's influence on the global estimate, concentrating privacy resources on the most informative time-steps. We theoretically prove that APBA satisfies each agent's global privacy budget under sequential composition while bounding estimation error. Experiments on real world sensor stream data demonstrate the efficiency and adaptability of APBA.

KEYWORDS

Multi-Agent Systems; Streaming Data; Differential Privacy; Adaptive Privacy Budget Allocation

ACM Reference Format:
Puspanjali Ghoshal and Ashok Singh Sairam. 2026. Utility Aware Adaptive Privacy Budget Allocation for Streaming Multi-Agent Systems. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 9 pages. https://doi.org/10.65109/DZYF2577

1 INTRODUCTION

The Multi-Agent System (MAS) paradigm is being increasingly deployed in domains like traffic coordination, distributed environmental monitoring and collaborative robotic teams. Each agent generates its own data stream and the fusion center analyzes these sequences to detect patterns. Time series data streams inadvertently leak sensitive and strategic information [1]. It can also leak information about the behavioral pattern of the agent [16].

Consider a network of smart meters in a neighborhood reporting real time electricity consumption to a utility provider. Usage patterns can reveal occupancy, appliance usage, or occupant schedules. However, these data are indispensable for grid balancing, demand forecasting, and fault detection. If too much noise is added uniformly across all reports, the utility provider will miss critical load surges or anomalies, degrading system reliability. Conversely, too little noise results in privacy risks getting compounded over time as temporal correlations can be exploited to reconstruct sensitive usage patterns. This example illustrates that for streaming MAS, privacy mechanisms must adapt to signal dynamics to avoid wasting budget in uninformative periods and thereby spending appropriately in high impact ones.

Differential Privacy (DP) has emerged as the standard for limiting inference risks by bounding the impact of a single datapoint on released outputs [15]. However, applying DP in streaming MAS introduces a critical tradeoff between privacy and utility. Excessive noise undermines situational awareness and control, whereas insufficient noise results in inference attack vulnerabilities.

DP budgets are finite and accumulate under sequential composition. Uniformly allocating budget or adding fixed noise per time-step leads to two key problems:

(i) Utility degradation: Informative time-steps with high variance are overly perturbed which affects tasks like forecasting or anomaly detection.
(ii) Inefficient budget use: Budgets are consumed at the same rate regardless of informativeness, leading to premature budget exhaustion or wasted budget during quiescent periods.

Thus, streaming MAS require privacy mechanisms that adaptively allocate noise over time, ensuring that privacy budgets are spent where they yield the highest utility. In this work, we propose *Adaptive Privacy Budget Allocation (APBA)*, a mechanism that dynamically distributes each agent's finite DP budget across time-steps. This allocation is guided by both the informativeness of the agent's observations and the potential impact of its data on the overall estimate of the system. Unlike fixed noise [10] or clipping based [17] approaches, APBA allocates more budget (less noise) during volatile periods and less budget (more noise) during stable periods, resulting in a principled privacy-utility tradeoff.

APBA is explicitly designed for analytics-first scenarios: its goal is to maximize estimation accuracy subject to a global privacy guarantee, not to maximize privacy at all costs. This design leads to a controlled but higher reconstruction similarity compared to

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/DZYF2577

conservative baselines, an intentional choice to preserve critical signal fidelity.

Our key contributions are as follows:

(i) We introduce APBA, a streaming capable, adaptive privacy mechanism for multi-agent data aggregation that prioritizes utility while respecting formal DP budgets.
(ii) We provide a theoretical proof that APBA satisfies each agent's total ϵ-DP budget under sequential composition and show analytically that it minimizes estimation variance compared to uniform allocation.
(iii) We empirically evaluate APBA on two real world sensor datasets (Intel Lab and NREL Wind Turbine) and demonstrate that it achieves a clear, tunable privacy-utility tradeoff.

2 RELATED WORK

Privacy preservation in MAS has received significant attention in recent years. In these settings, agents share data with a central fusion center. Sensitive information may be inferred from such reported values. Differential Privacy (DP) provides formal guarantees that individual contributions cannot be isolated and inferred [5]. However, traditional DP mechanisms often rely on fixed noise levels, which can degrade utility in heterogeneous or temporally evolving datasets [8, 14, 19].

Time series and streaming data introduce additional challenges due to temporal correlations. Huang et al. [12] combine local DP with a shuffling mechanism to amplify privacy and reduce noise accumulation. While effective in protecting individual measurements, this approach primarily uses static noise allocation and does not exploit temporal variations in signal uncertainty, limiting utility in streaming MAS. FedAPCA [18] addresses client heterogeneity by performing hierarchical aggregation with adaptive local DP across clients. However, it does not account for per time-step dynamics within individual agent signals.

Other adaptive DP mechanisms in federated learning adjust privacy or noise allocation over time. Kiani et al. [13] propose time adaptive DP for learning rounds, while Hong et al. [11] and Zhang et al. [20] dynamically adjust noise variance or gradient clipping during optimization. These methods demonstrate the benefits of adaptive allocation but focus on learning updates rather than raw streaming measurements. They also do not exploit per time-step signal uncertainty, which is critical in MAS.

Inference and reconstruction attacks further highlight the need for careful noise management [12]. In streaming MAS, static or poorly tuned noise can leave agents vulnerable, as adversaries can leverage temporal correlations across reports. In contrast, we propose Adaptive Privacy Budget Allocation (APBA), a mechanism that dynamically distributes each agent's total privacy budget across time-steps according to observed signal uncertainty and marginal contribution to the fusion center estimate. By adapting to temporal dynamics, APBA preserves high utility in volatile systems while providing resilience against reconstruction attacks in stable systems. This extends adaptive DP techniques to streaming multi-agent settings in a way prior methods do not.

3 SYSTEM MODEL

We consider a streaming MAS comprising of S autonomous agents. At discrete time-steps $t \in \{1, \ldots, T\}$, each agent i observes a vector $y_i(t) \in \mathbb{R}^n$ representing its local measurement. These observations are temporally correlated, reflecting the system dynamics or environment being monitored. Agents communicate their measurements to a central fusion center, which produces an estimate of an aggregated system parameter $\hat{\theta}(t)$, which is a statistic derived across all agents.

The following attributes are inherent to each agent i:

- Total privacy budget: $\epsilon_i^{\text{total}}$, representing the maximum allowable cumulative privacy loss over the streaming horizon.
- Remaining privacy budget: $\epsilon_i^{\text{remaining}}(t)$, which decreases over time as the agent perturbs observations.
- Local uncertainty measure: $u_i(t)$, computed using recent measurements to quantify variability or unpredictability. Suitable volatility metrics include variance and interquartile range among few.

At each time-step t, an agent perturbs its observation $y_i(t)$ using a Laplace mechanism with an adaptive privacy budget $\epsilon_i(t)$. Laplacian mechanism ensures pure $\epsilon_i(t)$-DP for data releases. It is well-suited for sequential streaming composition.

Definition 1. *Laplace Mechanism: Noise vector generated from the Laplacian distribution* $Lap\left(0, \frac{\Delta y}{\epsilon_i(t)}\right)$ *is added to the observation tuple. Here,* Δy *denotes the sensitivity of the observation, i.e., the maximum change in* $y_i(t)$ *resulting from any single agent's contribution.*

The magnitude of noise added is inversely proportional to the privacy budget $\epsilon_i(t)$. The adaptive allocation $\epsilon_i(t)$ is computed dynamically. Higher privacy budget $\epsilon_i(t)$ would imply less noise, while a lower value implies higher magnitude of added noise, and thereby higher levels of privacy.

The fusion center aggregates the perturbed reports $\{\tilde{y}_i(t)\}_{i=1}^{S}$ to compute an estimate of the system parameter $\hat{\theta}(t)$.

It continuously updates $\hat{\theta}(t)$ as new observations arrive. The adaptive privacy allocation directly affects estimation quality and enables the fusion center to receive informative reports while agents' privacy constraints are maintained.

The system operates continuously, with both agent observations $y_i(t)$ and adaptive privacy allocations $\epsilon_i(t)$ evolving over time. By exploiting temporal correlations, agents adjust noise magnitude according to signal volatility, ensuring the fusion center receives informative reports when needed. This dynamic allocation balances privacy and estimation utility, outperforming static, fixed noise mechanisms.

Each agent independently computes $\epsilon_i(t)$ and the fusion center aggregates these reports to estimate $\hat{\theta}(t)$. This distributed coordination is central to streaming MAS applications.

4 METHODOLOGY

We propose APBA, a streaming friendly mechanism that lets agents intelligently distribute differential privacy budgets over time while maximizing system utility. The proposed framework consists of three main components: (i) adaptive privacy budget allocation per

agent, (ii) local perturbation of observations, and (iii) fusion center aggregation.

At each time-step t, agent i allocates a fraction of its remaining privacy budget, $\epsilon_i(t)$, based on two main factors:

(1) Local signal uncertainty $u_i(t)$: Higher variance or volatility in $y_i(t)$ indicates that the observation carries more information. To preserve utility, more budget (less noise) is allocated when uncertainty is high.
(2) Influence on fusion center estimate $\Delta_i(t)$: Defined as

$$\Delta_i(t) = |\hat{\theta}(t) - \hat{\theta}_{-i}(t)|,$$

where $\hat{\theta}_{-i}(t)$ is the aggregated estimate without agent i. This quantifies the agent's contribution to the overall system estimate and guides allocation to maximize utility. Larger influence motivates higher budget allocation to ensure critical contributions are accurately captured. In practice, $\Delta_i(t)$ can be computed at the fusion center or estimated locally using historical reports.

The adaptive budget is computed as a normalized combination:

$$\epsilon_i(t) = \epsilon_i^{\text{remaining}}(t) \cdot \frac{\alpha u_i(t) + (1-\alpha)\Delta_i(t)}{\sum_{j=1}^{S} \left(\alpha u_j(t) + (1-\alpha)\Delta_j(t)\right)},$$

where $\alpha \in [0,1]$ balances the relative importance of local uncertainty versus global influence.

APBA allocates the per-timestep privacy budget based on an estimate of information importance derived from signal dynamics. Specifically, higher privacy budgets are assigned to timesteps where either (i) the local signal uncertainty $u_i(t)$ is large, or (ii) the agent's contribution to the global estimate, measured by $\Delta_i(t)$, is significant. Intuitively, these correspond to periods where accurate reporting is most valuable for downstream estimation.

Additionally, the total budget constraint is enforced:

$$\sum_{\tau=1}^{t} \epsilon_i(\tau) \le \epsilon_i^{\text{total}}.$$

The remaining privacy budget for agent i evolves according to:

$$\epsilon_i^{\text{remaining}}(t) = \epsilon_i^{\text{total}} - \sum_{\tau=1}^{t} \epsilon_i(\tau).$$

Once $\epsilon_i(t)$ is determined, agent i perturbs its observation using the Laplace mechanism:

$$\tilde{y}_i(t) = y_i(t) + \eta_i(t), \quad \eta_i(t) \sim \text{Laplace}\Big(0, \frac{\Delta y}{\epsilon_i(t)}\Big).$$

By sequential composition, the total privacy loss over T time-steps satisfies:

$$\sum_{t=1}^{T} \epsilon_i(t) \le \epsilon_i^{\text{total}},$$

ensuring that each agent adheres to its allocated ϵ-DP budget. The adaptive allocation concentrates budget on informative or high impact time-steps, improving utility while preserving formal privacy guarantees.

When $\epsilon_i^{\text{remaining}}(t) = 0$, the agent can no longer produce fully differentially private observations. Once the agent privacy budget is exhausted, one of the following approaches can be undertaken:

- Coarsened Reporting: Agents may provide aggregated temporal statistics requiring no additional budget.
- Budget Redistribution: Unused budgets from other agents can be reallocated to critical agents.
- Noise Masking: Agents may submit pure noise to maintain privacy while minimally contributing to system parameter estimation.

The algorithm 1 outlines our proposed method.

Algorithm 1 APBA for Streaming Multi-Agent Systems

Require: Number of agents S, total time-steps T, privacy budgets $\{\epsilon_i^{\text{total}}\}$, observations $\{y_i(t)\}$
Output: Streaming estimates $\{\hat{\theta}(t)\}$ and perturbed reports $\{\tilde{y}_i(t)\}$

Initialize $\epsilon_i^{\text{remaining}} \leftarrow \epsilon_i^{\text{total}}$ for all agents
for $t = 1$ to T **do**
  Compute local uncertainties $u_i(t)$ for all agents
  Compute fusion center impact $\Delta_i(t)$ for all agents
  for each agent i **do**
    if $\epsilon_i^{\text{remaining}} > 0$ **then**
      Allocate $\epsilon_i(t)$
      Generate perturbed report $\tilde{y}_i(t) = y_i(t) + \eta_i(t)$
      Update remaining budget: $\epsilon_i^{\text{remaining}} \leftarrow \epsilon_i^{\text{remaining}} - \epsilon_i(t)$
    else
      Apply coarsened reporting
      Set $\tilde{y}_i(t)$ accordingly
    end if
  end for
  Fusion center computes: $\hat{\theta}(t)$
end for
return $\{\hat{\theta}(t)\}, \{\tilde{y}_i(t)\}$

Both budget allocation and fusion center aggregation adapt dynamically over time. Agents assign larger budgets during periods of high volatility and smaller budgets during stable periods, achieving a principled tradeoff between privacy and system utility across the streaming horizon.

5 THEORETICAL GUARANTEES

We now provide formal guarantees for the proposed APBA mechanism. In Theorem 1, we show that APBA guarantees $((\epsilon_i^{\text{total}}, 0)$-differential privacy even under adaptive allocations. The guarantee continues to hold despite the gradual depletion of the privacy budget, while also minimizing estimation variance relative to uniform allocation. In Corollary 1, we show that the proposed approach preserves unbiasedness. Furthermore, the aggregated estimate converges asymptotically to the true parameter as S increases asymptotically, provided the estimation function at the fusion center is designed appropriately. The mechanism further degrades gracefully under partial agent dropout and achieves a strictly better privacy–utility tradeoff for streaming data characterized by high variance.

5.1 Differential Privacy Guarantees

We show that even with adaptive, nonuniform privacy budget allocations, the overall privacy loss is bounded.

Lemma 1 (Sequential Composition under Adaptive Allocation). *Let $\mathcal{M}_i(t)$ denote the mechanism applied by agent i in the time step t with budget $\epsilon_i(t)$. Then the cumulative mechanism over T time steps satisfies*

$$\mathcal{M}_i^{1:T} \text{ is } \Big(\sum_{t=1}^{T} \epsilon_i(t), 0\Big)\text{-DP.}$$

Proof. Each mechanism $\mathcal{M}_i(t)$ is $(\epsilon_i(t), 0)$-DP by construction. By the standard sequential composition theorem [6], applying these mechanisms over T time steps results in cumulative privacy loss equal to the sum of their individual budgets. Hence, $\mathcal{M}_i^{1:T}$ is $\big(\sum_{t=1}^{T} \epsilon_i(t), 0\big)$-DP. □

When $\epsilon_i^{\text{remaining}}(t) = 0$, APBA switches the agent i to coarse reporting. Let $\epsilon_i^{\text{coarse}}$ denote the (negligible) privacy cost of this fallback mechanism. The cumulative privacy loss is therefore the main concern.

$$\epsilon_i^{\text{cumulative}}(T) = \sum_{t=1}^{T} \epsilon_i(t) + \epsilon_i^{\text{coarse}} \le \epsilon_i^{\text{total}}.$$

Theorem 1 (Privacy Guarantee with Budget Exhaustion). *If each agent follows APBA with total budget ϵ_i^{total}, then APBA guarantees $(\epsilon_i^{total}, 0)$-differential privacy for the entire streaming horizon $t = 1, \ldots, T$, including after budget exhaustion.*

Proof. Lemma 1 ensures that the cumulative privacy loss is $\sum_{t=1}^{T} \epsilon_i(t)$. By construction, APBA enforces $\sum_{t=1}^{T} \epsilon_i(t) \le \epsilon_i^{\text{total}} - \epsilon_i^{\text{coarse}}$, and $\epsilon_i^{\text{coarse}} \ll \epsilon_i^{\text{total}}$. Hence, the total loss of privacy never exceeds $\epsilon_i^{\text{total}}$, which satisfies the global privacy guarantee. □

This formally establishes that each agent's cumulative privacy loss is bounded over the entire streaming horizon, ensuring no additional privacy loss occurs beyond budget exhaustion.

5.2 Utility Guarantees

We establish a bound on the estimation error incurred due to Laplace noise addition under APBA mechanism. This result quantifies how the local noise injections propagate through the fusion center estimator.

Theorem 2 (Bounded Estimation Error). *Let $\tilde{y}(t) = (\tilde{y}_1(t), \ldots, \tilde{y}_S(t))$ be the vector of reports from noisy agents. Let the fusion center estimator be any deterministic map $\hat{\theta}(t) = f(\tilde{y}(t))$, where $f : \mathbb{R}^S \to \mathbb{R}$ is $L(\ge 0)$-Lipschitz with respect to the Euclidean norm:*

$$|f(u) - f(v)| \le L\|u - v\|_2 \quad \forall\, u, v \in \mathbb{R}^N.$$

Then the mean squared estimation error satisfies

$$\mathbb{E}\big[\big|\hat{\theta}(t) - \theta(t)\big|^2\big] \le 2L^2 \sum_{i=1}^{S} \frac{(\Delta y)^2}{\big(\epsilon_i(t) + \epsilon_i^{\text{coarse}}\big)^2},$$

where $\theta(t) = f(y(t))$ is the estimator applied to the noiseless inputs.

Proof. Let $\eta = (\eta_1, \ldots, \eta_S)$. By definition, $\hat{\theta}(t) = f(y(t) + \eta)$ and $\theta(t) = f(y(t))$. By the Lipschitz property of f,

$$|\hat{\theta}(t) - \theta(t)|^2 = |f(y + \eta) - f(y)|^2 \le L^2\|\eta\|_2^2.$$

Taking expectations gives

$$\mathbb{E}\big[|\hat{\theta}(t) - \theta(t)|^2\big] \le L^2\mathbb{E}\big[\|\eta\|_2^2\big] = L^2 \sum_{i=1}^{S} \mathbb{E}[\eta_i^2].$$

For $b_i = \frac{\Delta y}{\epsilon_i(t) + \epsilon_i^{\text{coarse}}}$ and $\eta_i \sim \text{Laplace}(0, b_i)$ we have $\text{Var}(\eta_i) = 2b_i^2$, so

$$\mathbb{E}[\eta_i^2] = 2\left(\frac{\Delta y}{\epsilon_i(t) + \epsilon_i^{\text{coarse}}}\right)^2.$$

Substituting into the previous expression yields

$$\mathbb{E}\big[|\hat{\theta}(t) - \theta(t)|^2\big] \le 2L^2 \sum_{i=1}^{S} \frac{(\Delta y)^2}{\big(\epsilon_i(t) + \epsilon_i^{\text{coarse}}\big)^2},$$

which proves the claim. □

Corollary 1 (Convergence of Streaming Estimates). *If $y_i(t)$ are bounded and $S \to \infty$, then*

$$\lim_{S\to\infty} \mathbb{E}[\hat{\theta}(t)] = \theta(t),$$

up to noise variance, which vanishes at rate $1/S$ when $L = O(1/S)$. Thus, APBA preserves asymptotic unbiasedness of the aggregated estimate.

These results show that APBA not only controls the error induced by adaptive noise injection but also guarantees convergence of the aggregated estimate to the true parameter as the number of participating agents increases, providing a formal foundation for the mechanism's reliability and applicability to large-scale MAS.

5.3 Robustness to Dropout

We now extend the error analysis to account for agents that have exhausted their privacy budgets and thereafter switched to coarse reporting. The following corollary formalizes how such a scenario affects the mean squared estimation error.

Corollary 2 (Bounded Error under Agent Dropout). *Let $\mathcal{D} \subseteq \{1, \ldots, S\}$ be the subset of agents who have exhausted their privacy budgets and switch to coarse reporting. Let $\hat{\theta}(t) = f(\tilde{y}(t))$ be an unbiased L-Lipschitz estimator of $\theta(t) = f(y(t))$. Then the mean squared error satisfies*

$$\mathbb{E}\big[\big|\hat{\theta}(t) - \theta(t)\big|^2\big] \;\le\; 2L^2 \sum_{i \in \{1,\ldots,S\} \setminus \mathcal{D}} \frac{(\Delta y)^2}{\epsilon_i(t)^2} \;+\; \text{Var}_{\text{coarse}}(\mathcal{D}),$$

where $\text{Var}_{\text{coarse}}(\mathcal{D})$ denotes the total contribution to the error of the coarsened (fallback) reports of agents in $\mathcal{D}$.

Thus, the estimator exhibits graceful degradation as agents drop out, maintaining a finite error bound even when some agents use coarse reporting.

5.4 Adaptive Variance Reduction

We analyze the impact of adaptive budget allocation on the variance of the fusion center estimate. By adaptively directing the budgets, APBA reduces the noise contribution for informative time-steps. The following lemma formalizes this variance reduction relative to uniform per timestep allocation.

Lemma 2 (Variance Reduction via APBA). *Let $Var_{APBA}[\hat{\theta}(t)]$ and $Var_{Uniform}[\hat{\theta}(t)]$ be the variances under APBA and uniform allocation, respectively. Then*

$$Var_{APBA}[\hat{\theta}(t)] \leq Var_{Uniform}[\hat{\theta}(t)] \quad \textit{for high uncertainty time-steps.}$$

Proof. APBA assigns higher $\epsilon_i(t)$ to time steps with high $u_i(t)$ or $\Delta_i(t)$, reducing added noise precisely when the signal variability is highest. Uniform allocation cannot exploit this temporal heterogeneity, resulting in strictly higher variance during such periods. □

5.5 Privacy–Utility Optimality

Lemma 3 (Improved Privacy-Utility Tradeoff). *Let the total privacy budget of the i-th agent be ϵ_i^{total} over T timesteps. Define the effective weighted mean squared error as a heuristic bound:*

$$\text{MSE}_i := \sum_{t=1}^{T} \frac{w_t}{(\epsilon_i(t) + \epsilon_i^{\text{coarse}})^2}, \quad w_t \equiv w_i(t) := u_i(t)^2 + \Delta_i(t)^2,$$

with w_t being the timestep importance weight. Then the allocation that minimizes MSE_i under the budget constraint

$$\sum_{t=1}^{T} \epsilon_i(t) \leq \epsilon_i^{total}, \quad \epsilon_i(t) \geq 0$$

satisfies

$$\epsilon_i(t) + \epsilon_i^{\text{coarse}} \propto w_t^{1/3}.$$

Consequently, the APBA method reduces the expected error compared to uniform allocation in non-stationary settings.

Proof. In streaming MAS, some timesteps contribute more to the MSE due to larger signal uncertainty or larger influence on the fusion center. We want to minimize agent-specific MSE under the privacy budget, that is,

$$\min_{\epsilon_i(t)\geq 0} \sum_{t=1}^{T} \frac{w_t}{\left(\epsilon_i(t) + \epsilon_i^{\text{coarse}}\right)^2}, \quad \text{s.t.} \sum_{t=1}^{T} \epsilon_i(t) \leq \epsilon_i^{\text{total}}.$$

This is convex in $\epsilon_i(t)$. Introducing Lagrange multiplier $\lambda \geq 0$ and by subsequently setting the derivatives to zero, we have

$$\frac{\partial \mathcal{L}}{\partial \epsilon_i(t)} = -\frac{2w_t}{\left(\epsilon_i(t) + \epsilon_i^{\text{coarse}}\right)^3} + \lambda = 0$$

$$\implies \quad \epsilon_i(t) + \epsilon_i^{\text{coarse}} = \left(\frac{2w_t}{\lambda}\right)^{1/3}.$$

The optimal allocation is cubic-root weighted in w_t. APBA preserves this monotonic structure by assigning larger budgets to larger w_t, thereby moving toward the optimal allocation and reducing the weighted MSE relative to uniform allocation in heterogeneous streaming environments. APBA's adaptive rule therefore achieves an improved privacy-utility trade-off over uniform allocation which distributes the budget equally across all timesteps and therefore adds relatively more noise to high-impact steps, resulting in higher overall estimation error. Consequently, for the same privacy loss, APBA improves utility relative to uniform allocation, yielding a better privacy-utility tradeoff in non-stationary streaming settings. □

Although Theorem 2 provides a uniform upper bound, weighting timesteps reflects their true contribution to the expected squared error. This weighted MSE is convex and thus tractable for optimization, guiding APBA to allocate more budget where it matters most.

The update follows from minimizing a convex inverse-square error proxy under a linear global privacy constraint via Lagrangian relaxation. The resulting allocation scales as a power of the timestep importance, ensuring nonnegativity and satisfaction of the total privacy budget constraint.

6 EXPERIMENTATION

We evaluate our proposed APBA mechanism on multiple streaming multi-agent datasets, comparing it against state of the art baselines in terms of both utility and privacy leakage.

6.1 Datasets

We used two streaming datasets with distinct temporal characteristics:

(1) **Intel Lab Sensor Dataset [2]:**
- $S = 10$ sensors, $T = 1000$ time-steps, temperature and humidity streams.
- High-frequency measurements with occasional abrupt spikes due to sensor noise.

(2) **NREL Wind Turbine Dataset [7]:**
- $S = 8$ sites, $T = 1000$ time-steps, wind energy generation profiles.
- Smooth diurnal patterns with moderate day-to-day variability.

These two datasets allow us to evaluate APBA in both highly volatile and smoother temporal settings.

6.2 Baselines

The proposed approach is compared with two baseline approaches.

- **Shuffled DP** (S-DP) [12]: DP with random shuffling to reduce noise. Local ϵ measures the privacy of each user's locally randomized message, while central ϵ captures the overall privacy guarantee after shuffling and aggregating all messages.
- **FedAPCA** (F-APCA) [18]: Hierarchical aggregation with adaptive local DP, adjusting privacy budgets across clients to improve the privacy-utility tradeoff. The key parameter is the central differential privacy budget ϵ.
- **Variance Weighted** (VW): $\epsilon_i(t)$ is proportional to empirical variance. It is grounded in principles of distributed estimation, as in weighted consensus algorithms [3].

6.3 Adversarial Reconstruction Model

To evaluate the robustness of the proposed APBA mechanism, we consider a strong adversary performing a reconstruction attack on

reported noisy data streams. The adversary is assumed to have complete knowledge of the aggregation protocol, the noise distribution, and the global statistics of the system.

The adversary seeks to reconstruct the original signals $\{y_i(t)\}$ corresponding to the perturbed signal $\{\tilde{y}_i(t)\}$ of agent i. This is done by solving a convex optimization problem that balances fidelity to observed reports with temporal smoothness constraints. Specifically, the adversary solves the following:

$$\min_{\{\hat{y}_i(t)\}} \sum_{i=1}^{S} \sum_{t=1}^{T} \left(\hat{y}_i(t) - \tilde{y}_i(t)\right)^2 + \lambda \sum_{i=1}^{S} \sum_{t=2}^{T} \left(\hat{y}_i(t) - \hat{y}_i(t-1)\right)^2,$$

where $\hat{y}_i(t)$ denotes the reconstructed signal and $\lambda > 0$ is a regularization parameter controlling the temporal smoothness. The first term enforces proximity to the noisy observations, while the second term penalizes rapid fluctuations to exploit the expected temporal correlation in streaming data.

This attack can be efficiently solved using convex optimization libraries such as CVXPY [4]; which ensures that the solution converges to the global optimum. The reconstructed signals $\hat{y}_i(t)$ are then used to compute various metrics that quantify utility and privacy. This provides a rigorous, adversary-aware evaluation of privacy leakage under APBA and baseline mechanisms.

6.4 Evaluation Metrics

We use the following metrics for the quantification of utility and privacy.

- Mean Squared Error (MSE) and Mean Absolute Error (MAE) for utility.
- Breach count, average displacement, and resemblance [9] in a CVXPY reconstruction attack.

The trade-off between cumulative privacy loss and estimation error has been shown using Pareto curves.

7 RESULTS

7.1 APBA Epsilon Allocation Across time-steps

We first analyze how the proposed APBA mechanism allocates privacy budget across time steps. Figure 1 shows the resulting allocation profiles for two real-world datasets.

By construction, APBA increases per step $\epsilon_i(t)$ whenever the relative weight of the agent increases faster than the shrinkage of its remaining budget. This condition binds $\epsilon_i(t)$ to temporal variations in local uncertainty $u_i(t)$ and influence $\Delta_i(t)$ on the parameter being estimated.

The Intel dataset exhibits a strongly non-stationary pattern with transient spikes in measurement variance. As a result, $u_i(t)$ and $\Delta_i(t)$ grow in later time steps, driving the relative weight of the agent at the t-th timestep $w_i(t)/W(t)$ upwards despite budget depletion. Accordingly, $\epsilon_i(t)$ shows a gradual rise followed by a sharp end-of-horizon spike, as APBA concentrates the remaining budget on the most informative time steps. This adaptivity reduces noise during rapid state changes, improving the fidelity of the aggregate estimate.

In contrast, the NREL dataset is nearly stationary, with constant variance and influence over time. Here, $w_i(t)/W(t) \approx 1/S$ for all t,

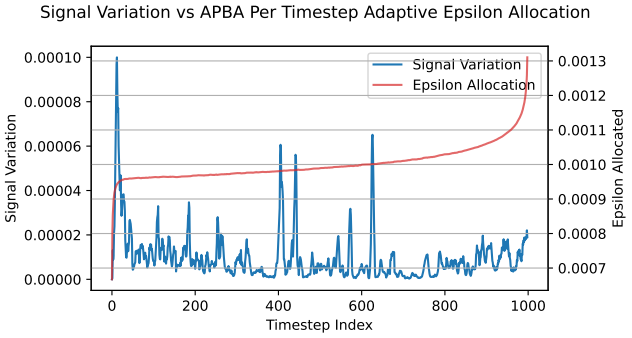

(a) Intel Sensor Dataset

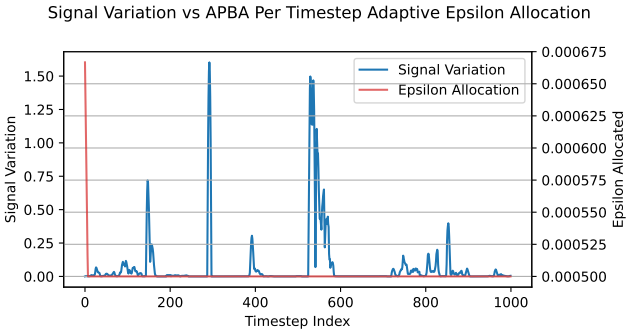

(b) NREL Wind Turbine Dataset

Figure 1: Epsilon allocation per time-step.

reducing the APBA to approximately uniform budget allocation:

$$\epsilon_i(t) \approx \frac{\epsilon_i^{\text{total}}}{T}, \quad \forall t. \tag{1}$$

The resulting flat profile in Fig. 1b confirms that APBA does not over allocate budget in stable periods, conserving resources while maintaining consistent privacy protection.

These findings validate APBA's design, it automatically spends more budget in high impact, high uncertainty regimes, and reverts to uniform spending when the signal is stationary. This adaptivity is crucial for balancing accuracy with differential privacy guarantees under sequential composition.

7.2 Estimation Accuracy

Next, we examine how adaptive budget allocation translates into estimation performance. Figure 2 compares the absolute error of APBA with two baselines: (i) Shuffled DP, which applies fixed per step ϵ followed by a shuffler, (ii) FedAPCA, a centralized mechanism with fixed ϵ applied to aggregated reports, and (iii) variance weighted baseline, where the privacy budget is allocated according to the empirical variance of the data.

For the Intel dataset (Fig. 2a), APBA achieves the lowest error in all time steps, reducing cumulative MAE by more than 90% compared to Shuffled DP and by roughly the same margin compared to FedAPCA. This gain comes from the APBA strategy of allocating more budget during periods of high variance (Fig. 1a), thus injecting less noise when observations are most informative. The Shuffled DP

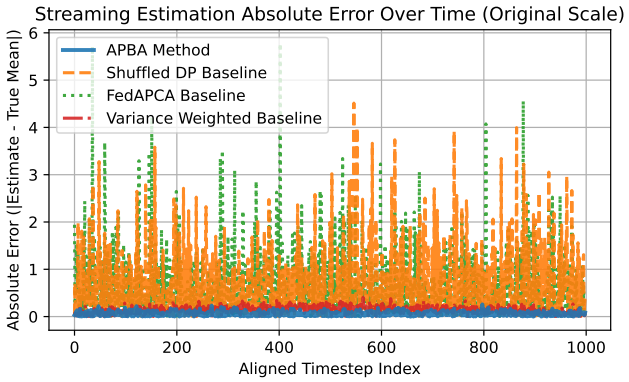

(a) Intel Sensor Dataset

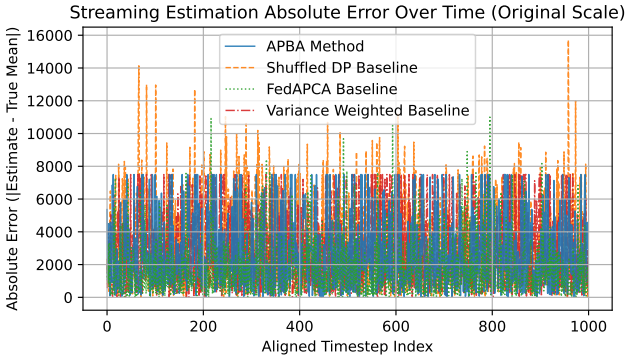

(b) NREL Wind Turbine Dataset

Figure 2: Absolute error per time-step.

baseline suffers from constant noise injection, leading to frequent large deviations from the true mean. FedAPCA performs better than Shuffled DP but remains less accurate, as its centralized noise addition does not exploit local signal heterogeneity.

For the NREL dataset (Fig. 2b), the performance gap is smaller but still positive. APBA reduces MAE by approximately 12% relative to shuffled DP. Since the dataset is stationary, APBA's privacy allocation becomes nearly uniform (Fig. 1b). Thus, the benefit of adaptivity is less pronounced but remains non-negative. These results confirm that APBA delivers consistent utility guarantees in varying signal dynamics.

The performance of our proposed method is comparable to the variance weighted baseline. However, unlike our proposed approach, the VW baseline does not offer the cumulative privacy guarantees. This gives our method an edge over the baseline and more suited for streaming MAS models.

7.3 Result Summary

Tables 1 and 2 summarize performance across all metrics.

In the Intel data set, APBA reduces MSE by over **99%** relative to S-DP and F-APCA baselines and its performance is comparable to VW baseline, confirming that adaptive allocation can achieve near perfect tracking under modest noise levels. The higher breach count of APBA and VW reflects an expected privacy-utility trade-off. More

Metric	S-DP [12]	F-APCA [18]	VW [3]	**APBA**
MSE	2.10426	2.18741	0.01141	**0.01206**
MAE	1.02807	1.04010	0.08612	**0.08903**
Breach Count	0.156	0.111	0.552	**0.529**
Displacement	9.13322	9.83289	2.04582	**2.14398**
Resemblance	0.50875	0.48366	0.51345	**0.51498**

Table 1: Results on Intel Sensor Dataset

Metric	S-DP [12]	F-APCA [18]	VW [3]	**APBA**
MSE	7.47687	2.14059	6.3606	**5.8088**
MAE	2.12446	1.02235	1.9487	**1.8645**
Breach Count	0.066	0.107	0.074	**0.038**
Displacement	22.08927	13.20502	17.91282	**20.43107**
Resemblance	0.49919	0.51891	0.50712	**0.50081**

Table 2: Results on NREL Wind Turbine Dataset

accurate reconstructions are inherently closer to the true trajectory. Importantly, APBA remains within the formal privacy budget, so this does not violate the DP guarantees. This guarantee improves the applicability of our method to streaming MASs, where both cumulative privacy and utility is required over multiple timesteps.

The substantial reduction in MSE and MAE for the Intel dataset in case of APBA and VW is due to the volatile nature of the Intel data set. Unlike APBA, fixed-(ϵ) mechanisms inject the same amount of noise at every step, which obscures signal during transient spikes. However, this gain is not universal. On the more stationary NREL data set, improvements are modest ($\approx$ 12% MAE reduction relative to shuffled DP), confirming that the large gains reflect dataset specific conditions rather than unrealistic performance.

In the NREL data set, APBA delivers competitive utility: MAE is reduced by ~12% and ~9% relative to shuffled DP and variance weighted baselines, respectively and the breach count is reduced by ~42% and ~50% respectively. Although MSE is higher than FedAPCA (5.80 vs. 2.14), APBA maintains better privacy–utility balance by avoiding excessive budget consumption in stable periods and preserving temporal structure, as indicated by resemblance scores.

The overhead of APBA is minimal, requiring only lightweight per-step updates to uncertainty and influence weights.

7.4 Pareto Analysis

Finally, we examine the Pareto frontier between privacy and utility (Fig. 3).

For the volatile Intel dataset, APBA lies near the bottom-right of the frontier, indicating high utility (low MSE) but with higher cumulative privacy loss, an intentional choice to minimize noise during transient spikes. For the smoother NREL dataset, APBA shifts toward the left-middle of the curve, achieving a lower privacy loss with only a modest utility improvement. This dataset specific behavior demonstrates that APBA dynamically adjusts its

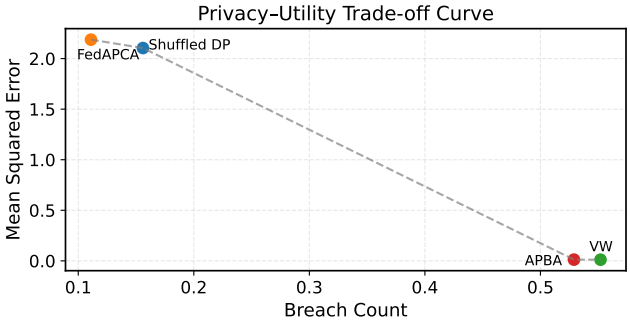

(a) Intel Sensor Dataset

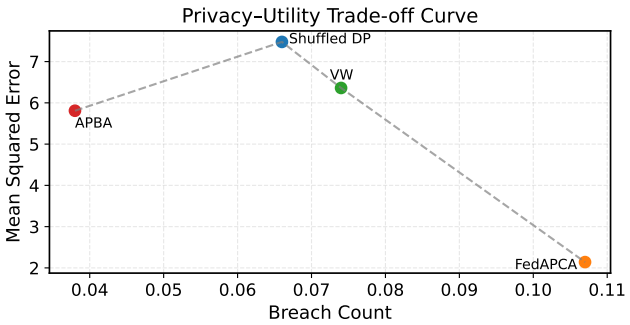

(b) NREL Wind Turbine Dataset

Figure 3: Privacy–utility Pareto frontier for Intel and NREL datasets.

position along the privacy–utility frontier, prioritizing accuracy when needed and conserving privacy budget otherwise.

Overall, these findings suggest that APBA provides a principled approach to balancing privacy and accuracy in cooperative multi-agent systems, making it a practical candidate for deployment in real world sensing and control applications.

8 CONCLUSION

We presented Adaptive Privacy Budget Allocation (APBA), a mechanism for dynamically distributing differential privacy budgets across time-steps in streaming MAS. Unlike static or client-level allocation schemes, APBA dynamically adjusts the per-step budget $\epsilon_i(t)$ based on local signal uncertainty and influence on the global estimate. Our theoretical analysis establishes that APBA guarantees the boundedness of the cumulative privacy loss of each agent. We assume a synchronized time model in which agents update privacy budgets concurrently at each round. The algorithm is fully decentralized in budget allocation, so increasing the number of agents increases computation and communication only proportionally. The decentralized APBA algorithm has $O(1)$ computation per agent per timestep, since each agent updates its budget using only local statistics (e.g., running variance and remaining budget). Overall system complexity is $O(S)$ per timestep for S agents, dominated by report aggregation at the fusion center. Since no pairwise coordination or cross-agent optimization is required, APBA is suitable for large-scale multi-agent sensing systems.

Preliminary analysis shows that APBA's performance is stable across variations in uncertainty weighting and budget scaling which implies robustness without extensive hyperparameter tuning. Empirical results demonstrate that APBA substantially improves estimation accuracy in nonstationary environments. In the Intel data set, APBA concentrates privacy budgets on informative time steps, reducing cumulative MAE by more than 90% relative to state-of-the-art baselines. In contrast, for the more stationary NREL dataset, APBA allocates budget almost uniformly over time, yielding smaller but consistent gains without early budget depletion.

This paper explores a synchronized model where the agents update the privacy budget concurrently. This method, however, can also be applied in asynchronous settings, which will be studied in future work. Future work will explore the incorporation of budget pacing to prevent premature depletion and the integration of predictive models to better anticipate uncertainty. We also plan to extend APBA to settings with partial observability or limited communication to enhance its robustness and applicability. Overall, APBA provides a theoretically grounded and practically effective framework for balancing privacy and utility in continuous-time cooperative sensing and control.

REFERENCES

[1] Tristan Allard, Hira Asghar, Gildas Avoine, Christophe Bobineau, Pierre Cauchois, Elisa Fromont, Anna Monreale, Francesca Naretto, Roberto Pellungrini, Francesca Pratesi, et al. 2024. Analyzing and explaining privacy risks on time series data: ongoing work and challenges. *ACM SIGKDD Explorations Newsletter* 26, 1 (2024), 49–58.

[2] Peter Bodik, Wei Hong, Carlos Guestrin, Sam Madden, Mark Paskin, and Romain Thibaux. 2004. Intel lab data. https://db.csail.mit.edu/labdata/labdata.html

[3] Erik Buchholz, Alsharif Abuadbba, Shuo Wang, Surya Nepal, and Salil Subhash Kanhere. 2022. Reconstruction Attack on Differential Private Trajectory Protection Mechanisms. In *Proceedings of the 38th Annual Computer Security Applications Conference* (Austin, TX, USA) *(ACSAC '22)*. Association for Computing Machinery, New York, NY, USA, 279–292. https://doi.org/10.1145/3564625.3564628

[4] Steven Diamond and Stephen Boyd. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* 17, 83 (2016), 1–5.

[5] Roel Dobbe, Ye Pu, Jingge Zhu, Kannan Ramchandran, and Claire Tomlin. 2018. Customized local differential privacy for multi-agent distributed optimization.

[6] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*. Springer, 265–284.

[7] Berk Erisen. 2018. Wind turbine scada dataset. *Kaggle* (2018).

[8] Puspanjali Ghoshal, Mohit Dhaka, and Ashok Sairam. 2024. On the effectiveness of differential privacy to continuous queries. *Service Oriented Computing and Applications* 18 (05 2024), 381–395. https://doi.org/10.1007/s11761-024-00397-9

[9] Puspanjali Ghoshal and Ashok Singh Sairam. 2024. Dimension Reduction via Random Projection for Privacy in Multi-Agent Systems. *arXiv preprint arXiv:2412.04031* (2024).

[10] Atefeh Gilani, Juan Felipe Gomez, Shahab Asoodeh, Flavio Calmon, Oliver Kosut, and Lalitha Sankar. 2025. Optimizing Noise Distributions for Differential Privacy. In *Proceedings of the 42nd International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 267)*, Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (Eds.). PMLR, 19505–19522. https://proceedings.mlr.press/v267/gilani25a.html

[11] Junyuan Hong, Zhangyang Wang, and Jiayu Zhou. 2022. Dynamic privacy budget allocation improves data efficiency of differentially private gradient descent. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 11–35.

[12] Chenxi Huang, Chaoyang Jiang, and Zhenghua Chen. 2023. Shuffled Differentially Private Federated Learning for Time Series Data Analytics. In *2023 IEEE 18th Conference on Industrial Electronics and Applications (ICIEA)*. IEEE, 1023–1028.

[13] Shahrzad Kiani, Nupur Kulkarni, Adam Dziedzic, Stark Draper, and Franziska Boenisch. 2025. Differentially private federated learning with time-adaptive privacy spending. *arXiv preprint arXiv:2502.18706* (2025).

[14] Ke Pan and Kaiyuan Feng. 2023. Differential privacy-enabled multi-party learning with dynamic privacy budget allocating strategy. *Electronics* 12, 3 (2023), 658.

[15] Rishabh Subramanian. 2023. Have the cake and eat it too: Differential Privacy enables privacy and precise analytics. *Journal of Big Data* 10, 1 (2023), 117.
[16] Nazanin Takbiri, Amir Houmansadr, Dennis L. Goeckel, and Hossein Pishro-Nik. 2017. Matching Anonymized and Obfuscated Time Series to Users' Profiles. *IEEE Transactions on Information Theory* 65 (2017), 724–741.
[17] Om Thakkar, Galen Andrew, and H. B. McMahan. 2019. Differentially Private Learning with Adaptive Clipping. In *Neural Information Processing Systems*. https://api.semanticscholar.org/CorpusID:150373827
[18] Jie Wang, Zhiju Zhang, Jing Tian, and Hongtao Li. 2024. Local differential privacy federated learning based on heterogeneous data multi-privacy mechanism. *Computer Networks* 254 (2024), 110822.
[19] Zhiqiang Wang, Xinyue Yu, Qianli Huang, and Yongguang Gong. 2024. An adaptive differential privacy method based on federated learning. *arXiv preprint arXiv:2408.08909* (2024).
[20] Jiaojiao Zhang, Dominik Fay, and Mikael Johansson. 2024. Dynamic privacy allocation for locally differentially private federated learning with composite objectives. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 9461–9465.