INDIAN INSTITUTE OF TECHNOLOGY GUWAHATI
PhD-17 SHORT ABSTRACT OF THESIS

Name of the Student	:	Puspanjali Ghoshal
Roll Number	:	206123105
Programme of Study	:	Ph.D.

Thesis Title: Preserving Inference Privacy in Multi-Agent Systems

Name of Thesis Supervisor(s)	:	Prof. Ashok Singh Sairam
Thesis Submitted to the Academic Division	:	Yes
Date of completion of Thesis Viva-Voce Exam	:	03/08/2026
Key words for description of Thesis Work	:	Inference Privacy, Multi-Agent Systems, Utility, Data Streams, Distributed Intelligence

SHORT ABSTRACT

A multi-agent system (MAS) involves autonomous agents observing and communicating with a fusion center, which aggregates the data to infer one or more system parameters with the highest possible accuracy. However, besides the system parameters, an honest-but-curious fusion center may infer secondary sensitive information or create a profile of an agent without its consent. These attacks constitute an inference privacy breach. This thesis aims to develop robust and efficient mechanisms to prevent such breaches in MASs. Since any entity other than the data owner can be a potential adversary, preserving inference privacy requires decentralized data sanitization. Moreover, the limited computational resources of autonomous agents necessitate lightweight sanitization mechanisms. As data sanitization inevitably modifies the raw observations and reduces utility, privacy mechanisms must preserve sufficient utility for accurate distributed inference at the fusion center.

The thesis investigates the applicability of differential privacy (DP) to mobile agents that continuously share their location information with a third party for location-based services. We show that the privacy guarantee provided by standard DP with fixed per-timestep budget allocation degrades over time due to sequential composition. Thus, we propose a perturbation generation mechanism that provides improved privacy preservation for continuous queries.

The thesis extends this framework to higher-dimensional data streams by proposing an adaptive privacy budget allocation scheme. The proposed method allocates the total available budget over timesteps according to the informativeness of the data and the importance of the agent to the system. Informative timesteps are prioritized to yield higher utility while maintaining cumulative privacy.

Next, the thesis moves on from perturbation-based methods to dimension reduction based methods. Exploiting the similarity between the MAS framework and robust concept deduction, a dynamic projection method is developed. The projection matrix evolves based on the previous raw observation and its sanitized counterpart, resulting in an adaptive data-dependent mechanism that achieves an effective privacy-utility tradeoff.

Finally, adaptive modulo hashing is introduced. The proposed method eliminates the requirement of a fixed divisor and produces data-dependent hash values for resource-constrained multi-agent systems. Collectively, the thesis proposes decentralized, lightweight, and computationally efficient solutions for preserving inference privacy in agent-based distributed systems.