Schedule Indistinguishability: Applying Differential Privacy Models to Protect Runtime Timing Behaviors

Maryam Ghorbanvirdi and Sibin Mohan

Department of Computer Science, The George Washington University

Email: {maryam.ghorbanvirdi, sibin.mohan}@gwu.edu

Abstract

Differential Privacy (DP) is a popular framework for data analysis that has evolved from a theoretical concept to practical applications over the last 20 years, protecting critical data in extensive deployments. Recent studies have examined DP's suitability for protecting real-time systems (RTS), where runtime data itself may be a target for privacy leakage. Previous research has investigated ϵ-differential privacy (ϵ-DP) based scheduling, that adds Laplace noise to execution times to reduce such hazards. However, Laplace noise is vulnerable to persistent adversarial attacks because of its predictability and set privacy budgets. Inspired by Rényi Differential Privacy (RDP), we propose a unique adaptive privacy-preserving scheduling technique in this study. Our method uses Gaussian-based noise injection guarantees schedule indistinguishability even when subjected to long-term adversarial monitoring. Additionally, we test our method in a client-server environment, where an adversary can either watch system replies from the server side or inject timing-based attacks from the client side. According to our experimental findings, RDP-based scheduling works better than ϵ-DP techniques in preserving indistinguishability and minimizing performance deterioration.