CERIAS

The Center for Education and Research in Information Assurance and Security

Forgetting the Forgotten with Letheia: Concealing Content Deletions in Presence of Persistent Observers

Mohsen Minaei[1], **Mainack Mandal**[2], **Patrick Loiseau**[2], **Krishna Gummadi**[2], **Aniket Kate**[1]

[1]Purdue University, USA [2]MPI-SWS, Germany

1. Data Enormity & Data Exposure

Today people freely open up about their personal life and opinions on online social platforms, causing *long-term exposure* to intended recipients and unintended data scavengers.

2. Privacy Issues and Deletions

Privacy Attacks.

- Harassment and blackmail of victims for their sensitive posts
- Given the enormity of accessible data, large scale mining may not be economical for the adversary

Data Deletions.

- Users request to withdraw their data are honored by several platforms
- *Streisand effect*: an attempt to hide some information has the unintended consequence of brings particular attention of public to it
- Data deletions make the data scavengers task significantly easy!

3. State of the Art and Corresponding Problems

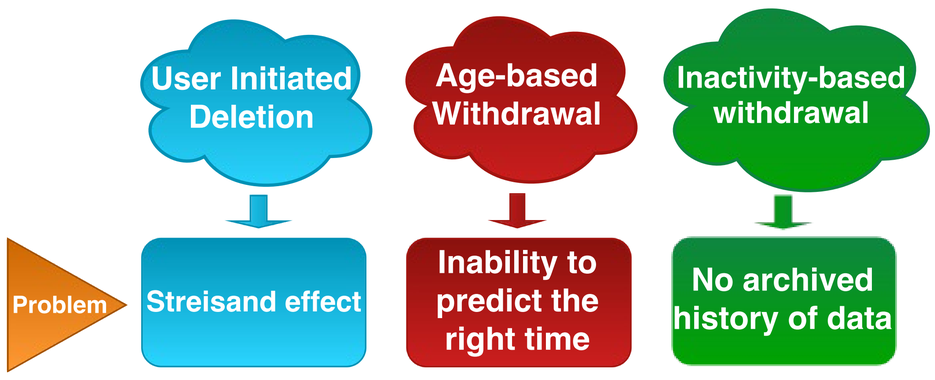

4. Letheia: Intermittent Withdrawals

Our proposal employs an *intermittent withdrawal mechanism* using two time distributions: an up distribution, and a down distribution.

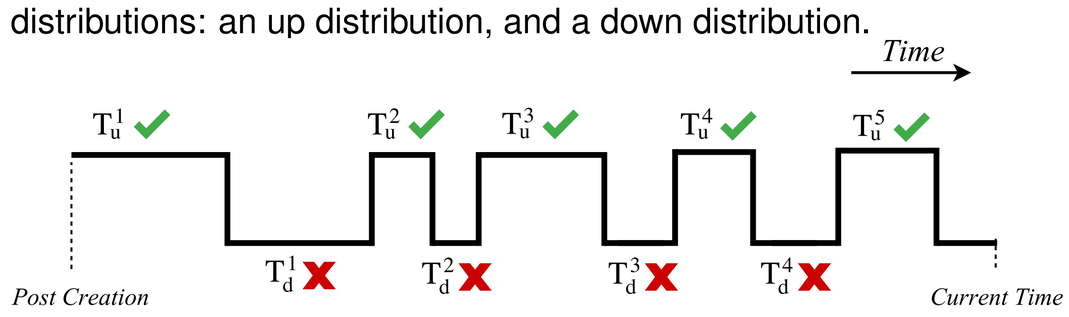

T_u^i is the duration of an up phase(post is visible to all) and T_d^i is the duration of a down phase(post is hidden to all). we toggle between the up and down durations as long as the post has not been deleted

5. Security Goals in Letheia

Deletion Privacy

- Uncertainty about a post being deleted or just temporarily withdrawn
- Defined as a log likelihood ratio (LLR) between a post being in a down duration vs. the post being deleted at a particular time

System Utility

- The overall percentage of posts being available for interactions
- Captured by the volume of interactions a post receives as time passes.

Adversarial Overhead

- # of undeleted posts flagged as deleted (false-positives) that the adversary has to investigate along with the detected deleted posts (true-positives)
- Captured by the *precision* measure, $\text{precision} = \frac{\text{true-positives}}{\text{true-positives} + \text{false-positives}}$

7. Experimental Results

- Evaluated the effectiveness of Letheia with interaction data from Twitter.
- Showing the trade-off between privacy, utility and the adversarial overhead when we chose different mean down durations for negative binomial distribution.

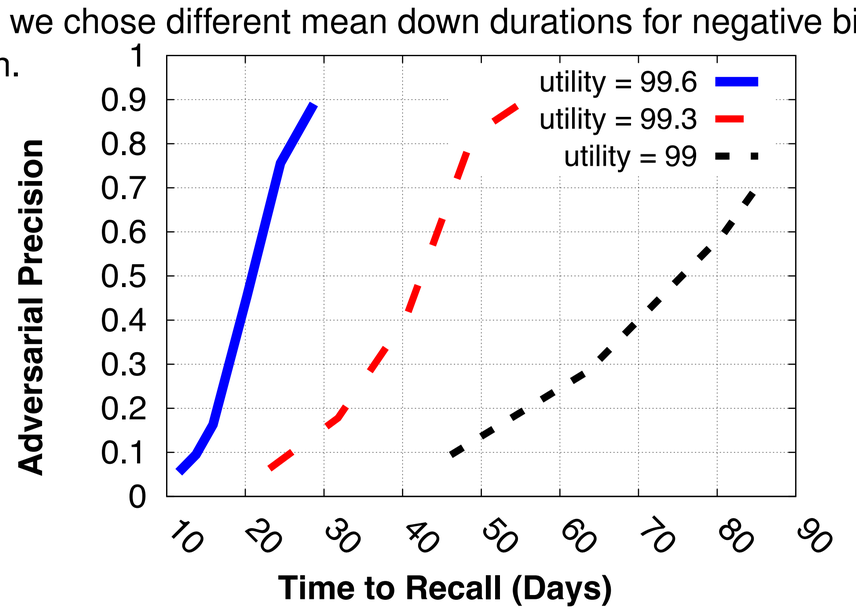

Time to recall: adversary's wait duration to make a decision about a post being hidden or deleted. Each time to recall is associated with a LLR value.

6. Distribution Selection

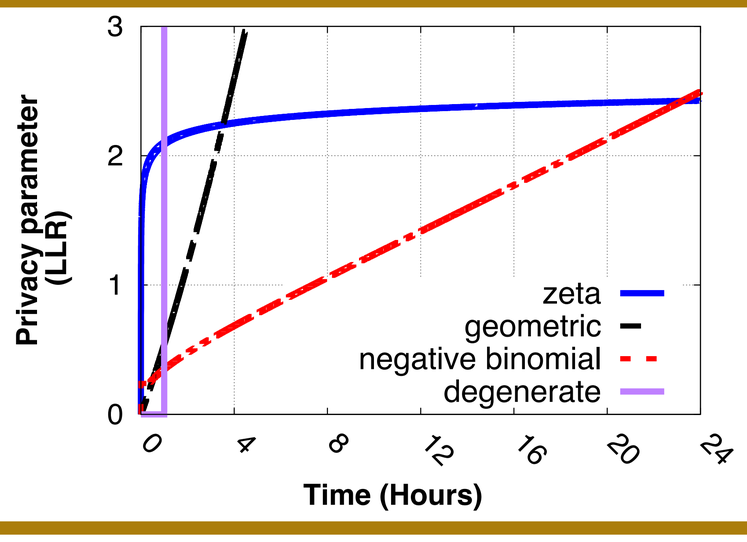

Variation of privacy parameter (LLR) with time for four choices for down distributions. Increase in LLR signifies decrease in privacy.

8. Acknowledgment

This project has been supported by the CERIAS/Intel Research Assistantship.