Privacy-Preserving Distributed Storage and Computation*

Michael Goodrich
U. California, Irvine

Michael Mitzenmacher
Harvard U.

Roberto Tamassia
Brown U.

This project aims at developing efficient methods for protecting the privacy and integrity of computations on outsourced data in distributed settings.

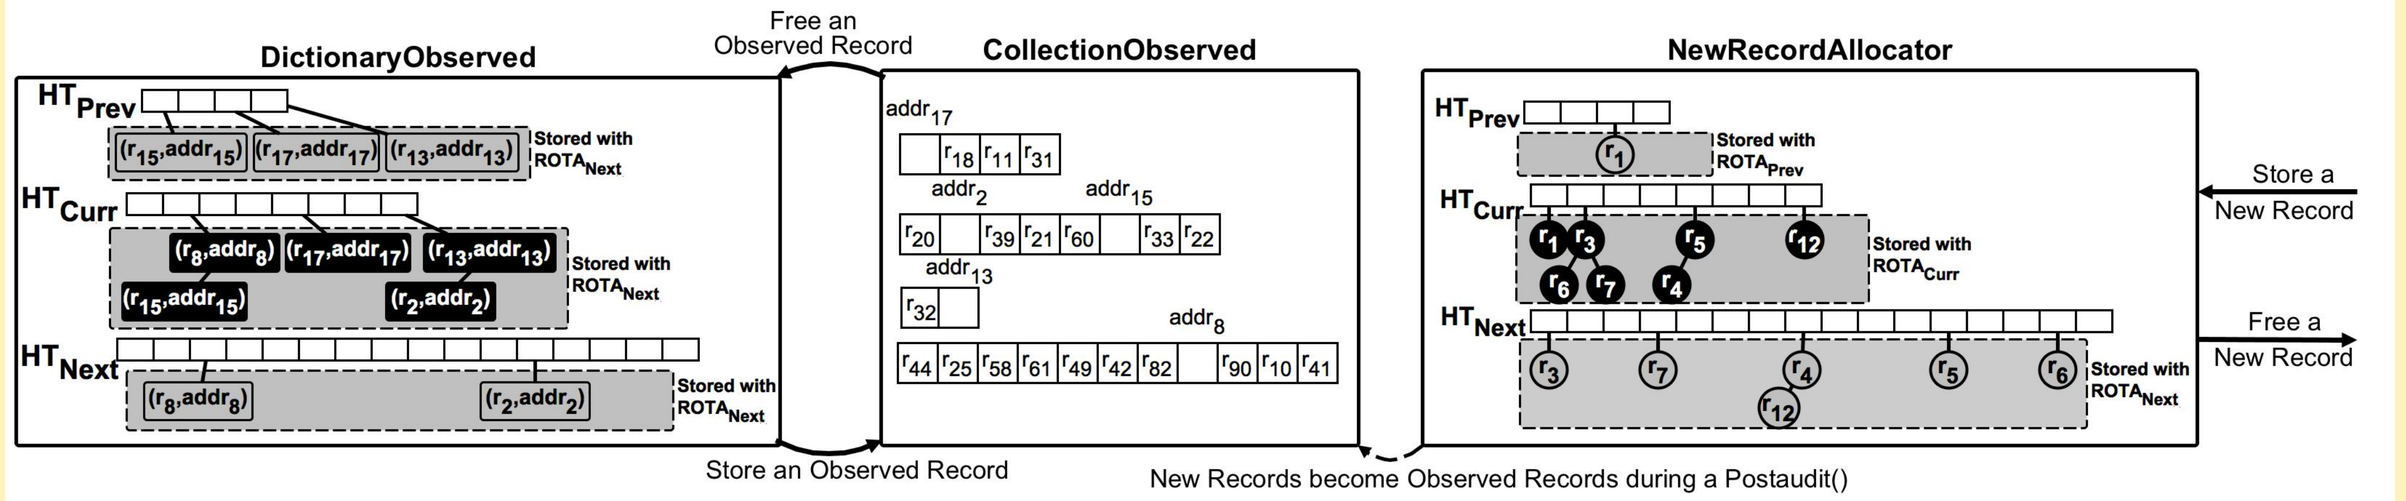

Models

- Cloud computing
- Two-party model
- Three-party model

Integrity

- Verifiable data structures
- Accumulators

Privacy

- Zero-knowledge
- History independence
- Watermarking

Auditable Data Structures

- A new model of history independence
- Snapshot of data structure taken at arbitrary times
- Awareness of and reaction to audit
- Auditor learns nothing about history of operations leading to current state
- Generic and efficient memory manager
- Applications include voting machines

History Independent Hash Tables

- Based on linear probing
- Secure against memory snapshots attacks
- Secure against collision timing attacks
- Insert/delete up to 2x faster than previous work
- Find up to 2x faster than previous work

ZK Verifiable Data Structures

- Owner outsources database to server and periodically updates it
- Client queries and server returns answer and proof
- Proof does not reveal anything about database beyond current and previous answers

Graph Watermarking

- General graph watermarking framework
- Information encoded in small changes to edges
- Resiliency to adversarial modifications
- Simple and secure edge-flipping marking schemes for Erdős-Rényi power-law random graphs

*NSF grants 1228639, 1228598, 1228485.

National Science Foundation
WHERE DISCOVERIES BEGIN

The 3rd NSF Secure and Trustworthy Cyberspace
Principal Investigator Meeting
January 9-11, 2017
Arlington, Virginia