

### 1. Introduction

- **Geo-distributed Systems:** Data is replicated across regions. Systems usually assign data ownership to a single region to maintain consistency.
- **Remastering:** Data ownership dynamically moves closer to active users to increase throughput and reduce latency.
- **Privacy Constraints:** Transactions are restricted by privacy regulations (e.g., GDPR).
- **The Gap:** Existing systems optimize for performance, largely ignoring privacy constraints in data movement decisions.

### Research Question

What is the impact of privacy-constrained dynamic data movement on the throughput and latency of transactions in Detock?

### 2. Background & Implementation

- **Remastering Strategies:** Approaches range from simple policies (SLOG/PNUTS, LEAP) to more advanced policies (DynaMast, Tuba).
- **Privacy Laws:** Privacy Laws such as GDPR, FADP, and CCPA legally restrict data movement across regions.
- **Implementation:** Our implementation allows each row to be marked as non-movable. The remaster controller will then block remasters based on this label.

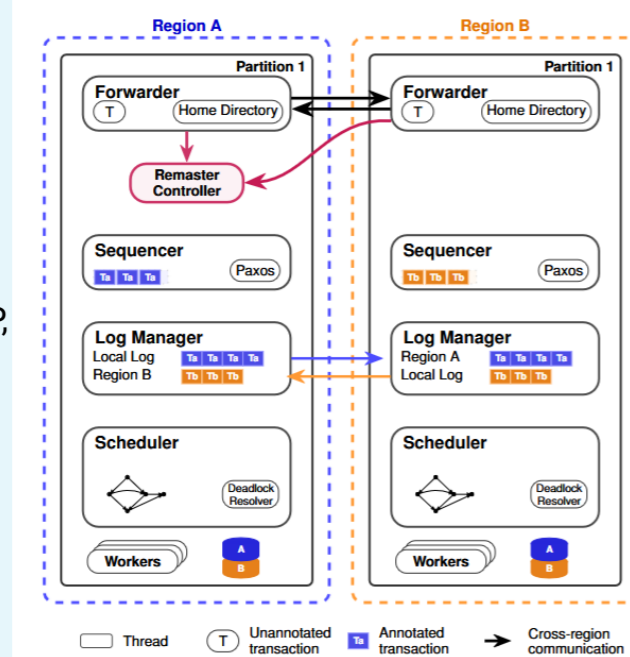


Figure: The Detock architecture with the remaster controller.

### 3. Experimental Setup and Results

- Evaluated performance across varying percentages of private rows using three distinct benchmarks: YCSB, DeathStar Movie, and TPC-C
- Compared SLOG-inspired, DynaMast-inspired, and co-access-based remastering strategies.
- Each benchmark used specific workloads made to encourage remastering.
- Collected metrics such as throughput, latency, remastering rate, and cost.

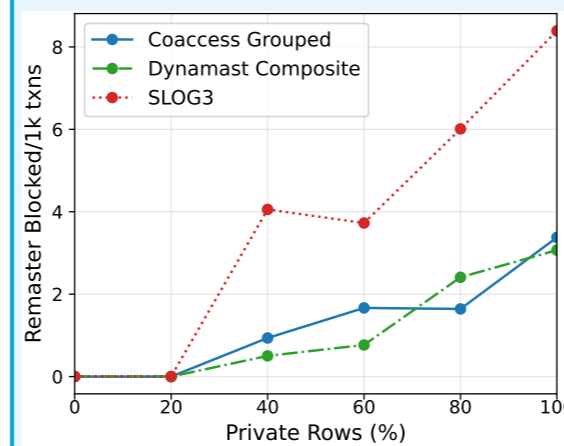

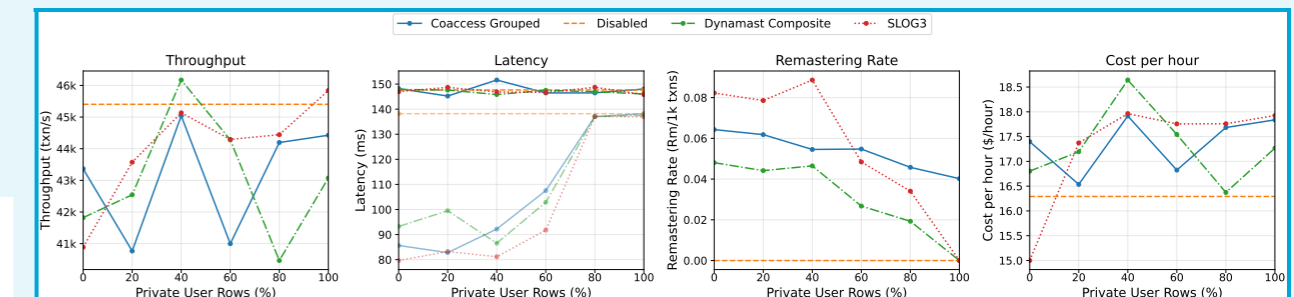


Figure: Remasters blocked due to privacy under the TPC-C workload.

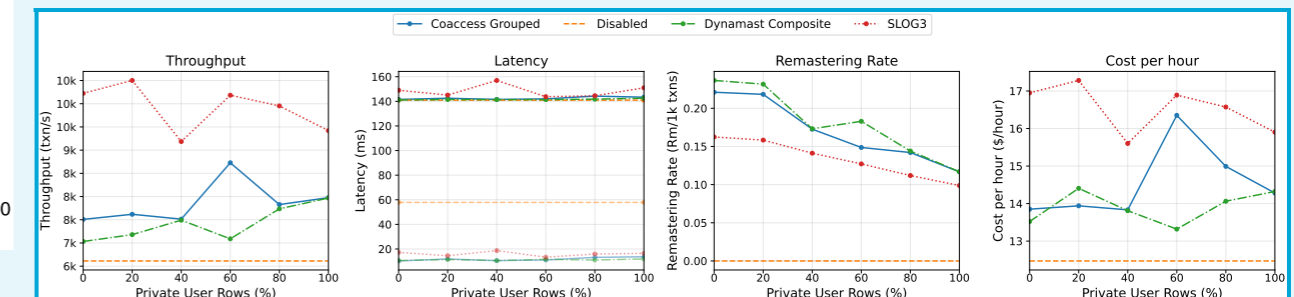
#### YCSB workload



#### Movie workload



#### TPC-C workload



### 4. Discussion

#### Impact of Privacy Constraints:

- Remastering still stays effective under moderate privacy levels.
- Performance drops happen once a large fractions of rows become immovable.

#### Limitations & Future Work:

- Used workloads encouraging remastering to study privacy effects; design realistic benchmarks built with remastering and privacy in mind.
- Privacy model is simplified; extend Detock to comply with existing privacy laws.
- Design a specific remastering policy for privacy.

### 5. Conclusion

- This work extends Detock to support privacy constraints during remastering.
- The results demonstrate that moderate privacy constraints have minimal performance impact.
- Ultimately, privacy constraints should be considered for dynamic data movement.