

CaSA: End-to-end Quantitative Security Analysis of Randomly Mapped Caches

Thomas Bourgeat, Jules Drean, Yuheng Yang, Lillian Tsai, Joel Emer, Mengjia Yan

Published at MICRO'20

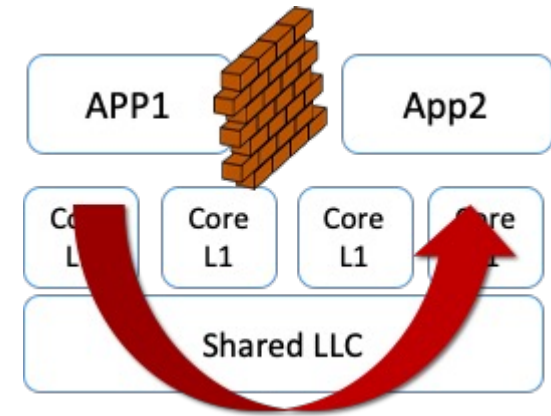


Problem: Incomplete Security Analysis


- Cache side channels are a serious security threat
- Promising mitigation: randomly mapped cache



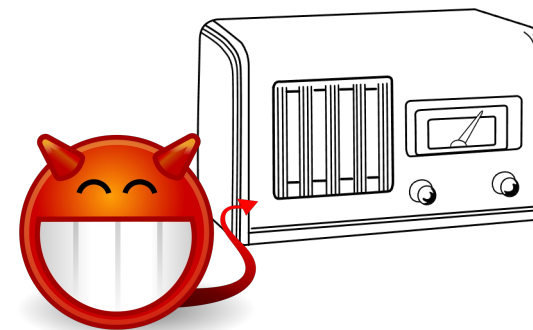
- The security property is not well understood



Key Insights: Telecommunication Analogy




Cache Channel



- Contributions:


- CaSA leverages concepts from telecommunications to enable quantitative analysis
- An end-to-end communication paradigm to enable comprehensive analysis
- New findings that refute common beliefs

Cache Side-channel Attacks



Attacker monitor victim's behavior through cache conflicts.

Cache Side-channel Attacks




Using black-box mapping function increases the cost to build an Eviction Set to:
 $O(N)$ (with N the number of lines in the cache)

Security Metrics

- Community intuition on Security Metrics:
 - “How hard it is to build an eviction set” is a good quantitative notion of security
- State-of-the-art secure cache design approaches:
 - Dynamic mapping
 - Non-deterministic mapping
- Our work:
 - This security metric can be misleading
 - Both design approaches fail to provide security

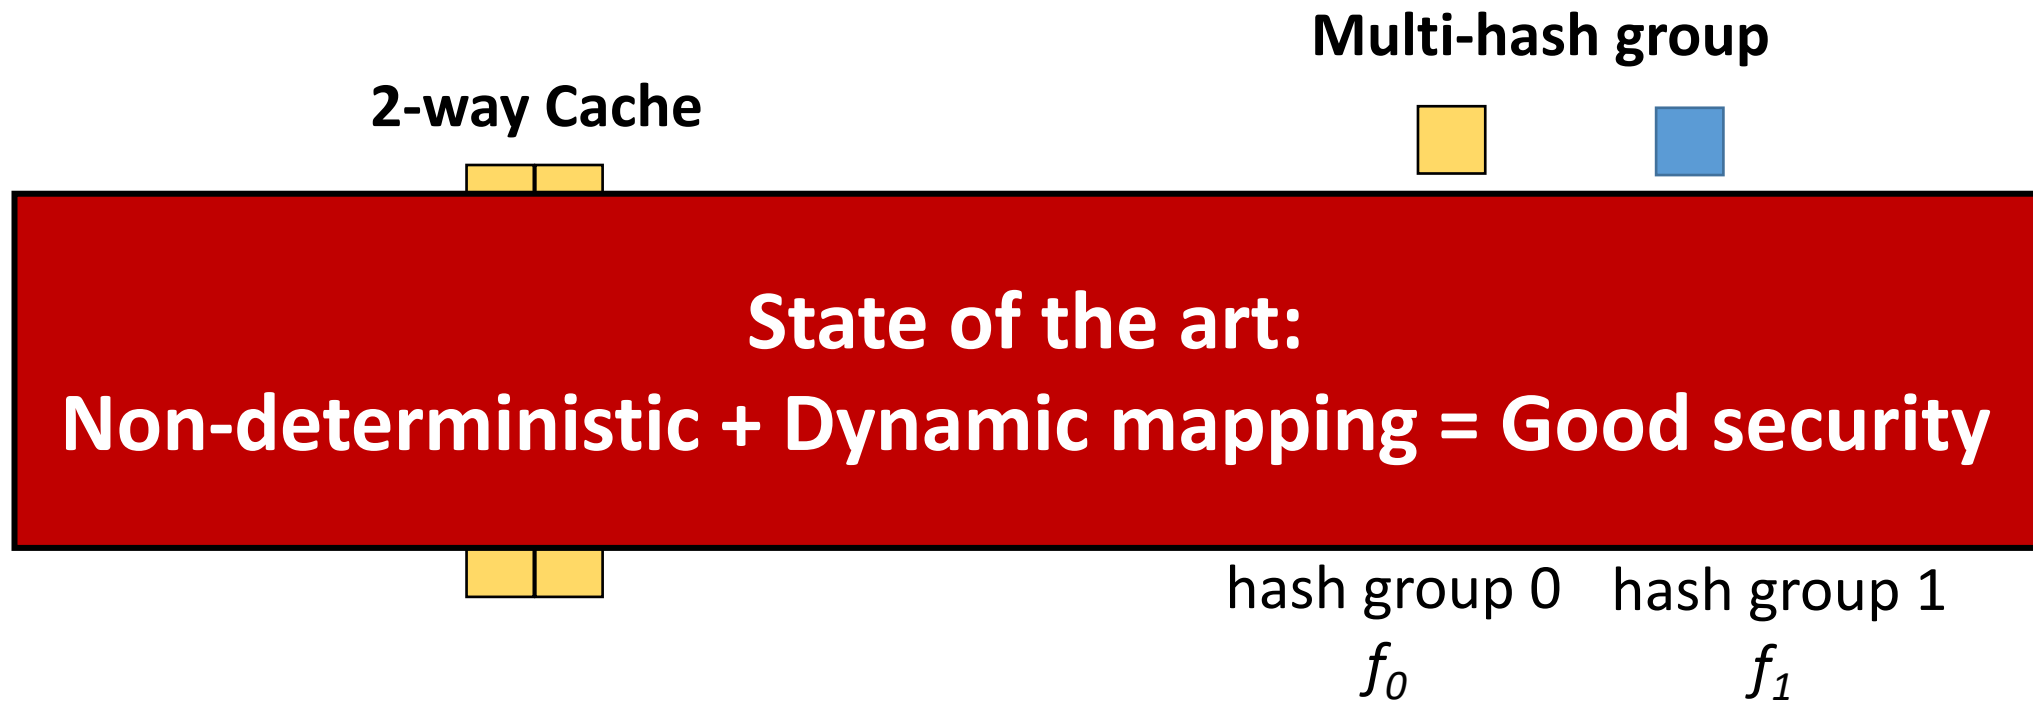
Dynamic Mappings

- Common belief: attacks can not happen across epochs
- Dynamic remapping incurs performance overhead




Non-Deterministic mapping

- Make conflict relationship between addresses non-deterministic




Traditional Analysis

- Hard-conflict addresses:
 - Guarantee eviction
 - Difficult to obtain




- Soft-conflict addresses:
 - Easy to obtain
 - Need many of such addresses to reliably evict addresses



- Narrowly focus on eviction set construction and lose the bigger picture.
- Only want to create a one-to-one map from micro-architecture events to secret

End-to-end Communication Paradigm

- Leverage the concepts from telecommunication
- Trade-off between calibration and signaling
 - Long time on calibration → shorter time needed for signaling
 - Short time on calibration → longer time needed for signaling



New Security Metric

~~“How difficult to construct an eviction set”~~

End-to-end communication cost in **Calibration + Signaling**



Statistical Representation of Signals

- Signal: a random variable “X”
 - Describes the number of misses observed by the attacker
 - Follow a probability distribution

- Example:

	Prob observing 0 miss	Prob observing 1 miss
Victim accesses	0.75	0.25
No victim accesses	1	0

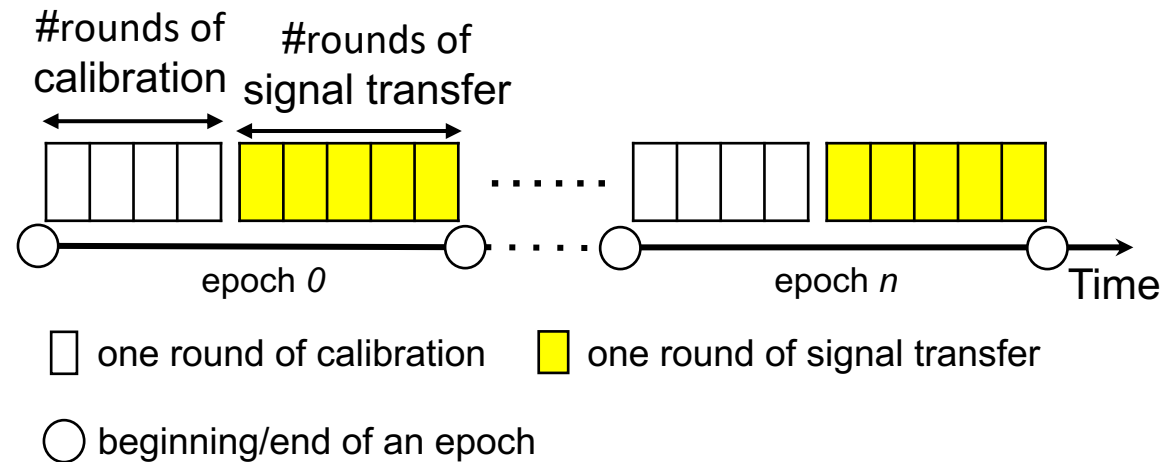
- Successfully covert the problem to a statistical analysis problem
 - How many samples are needed to distinguish the two distributions?

Two Insightful Findings

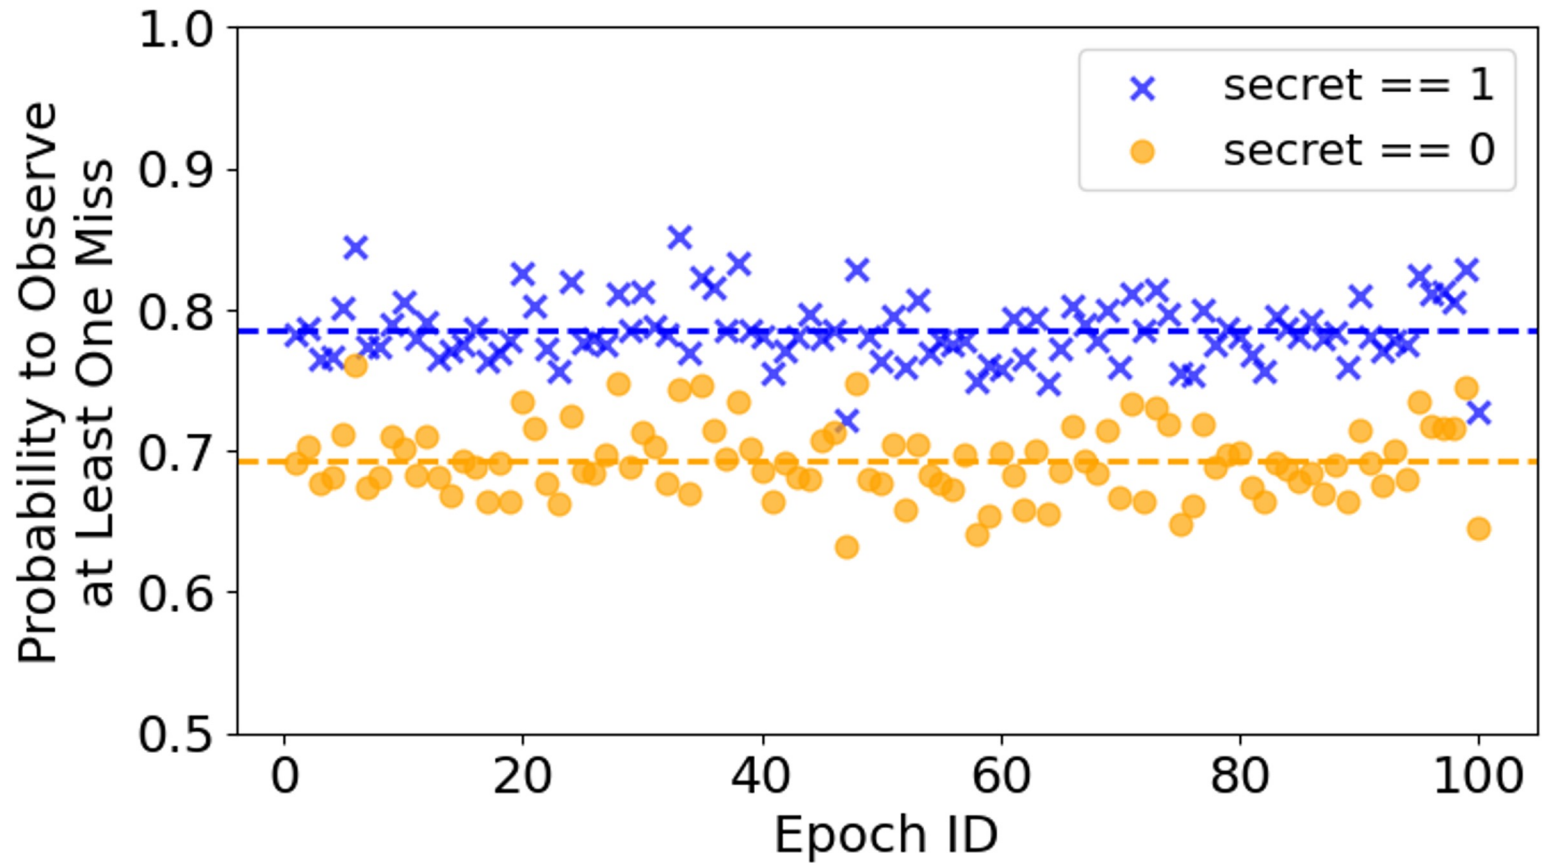
- Cross-epoch communication is possible
- Spending maximum resources on calibration is not the best strategy

Cross-epoch Communication

- In each epoch:

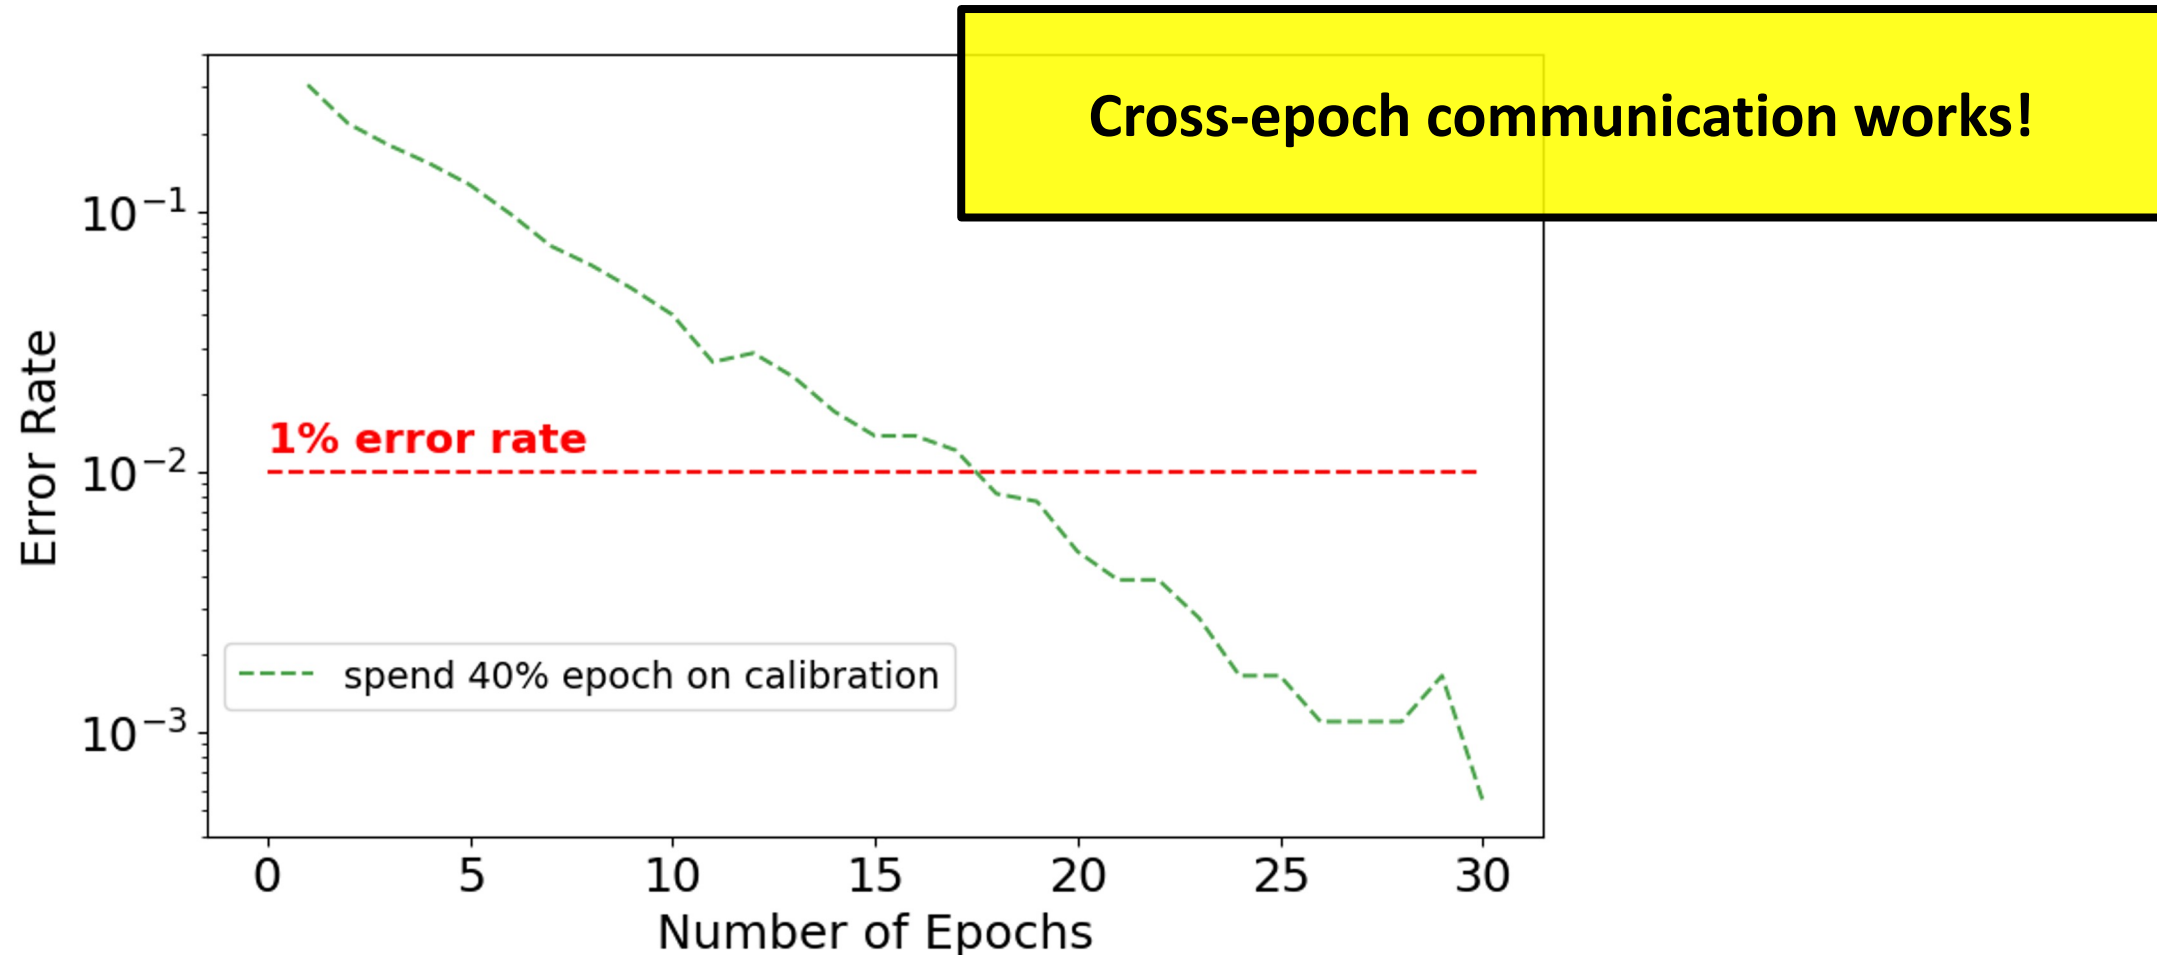


Cross-epoch Communication

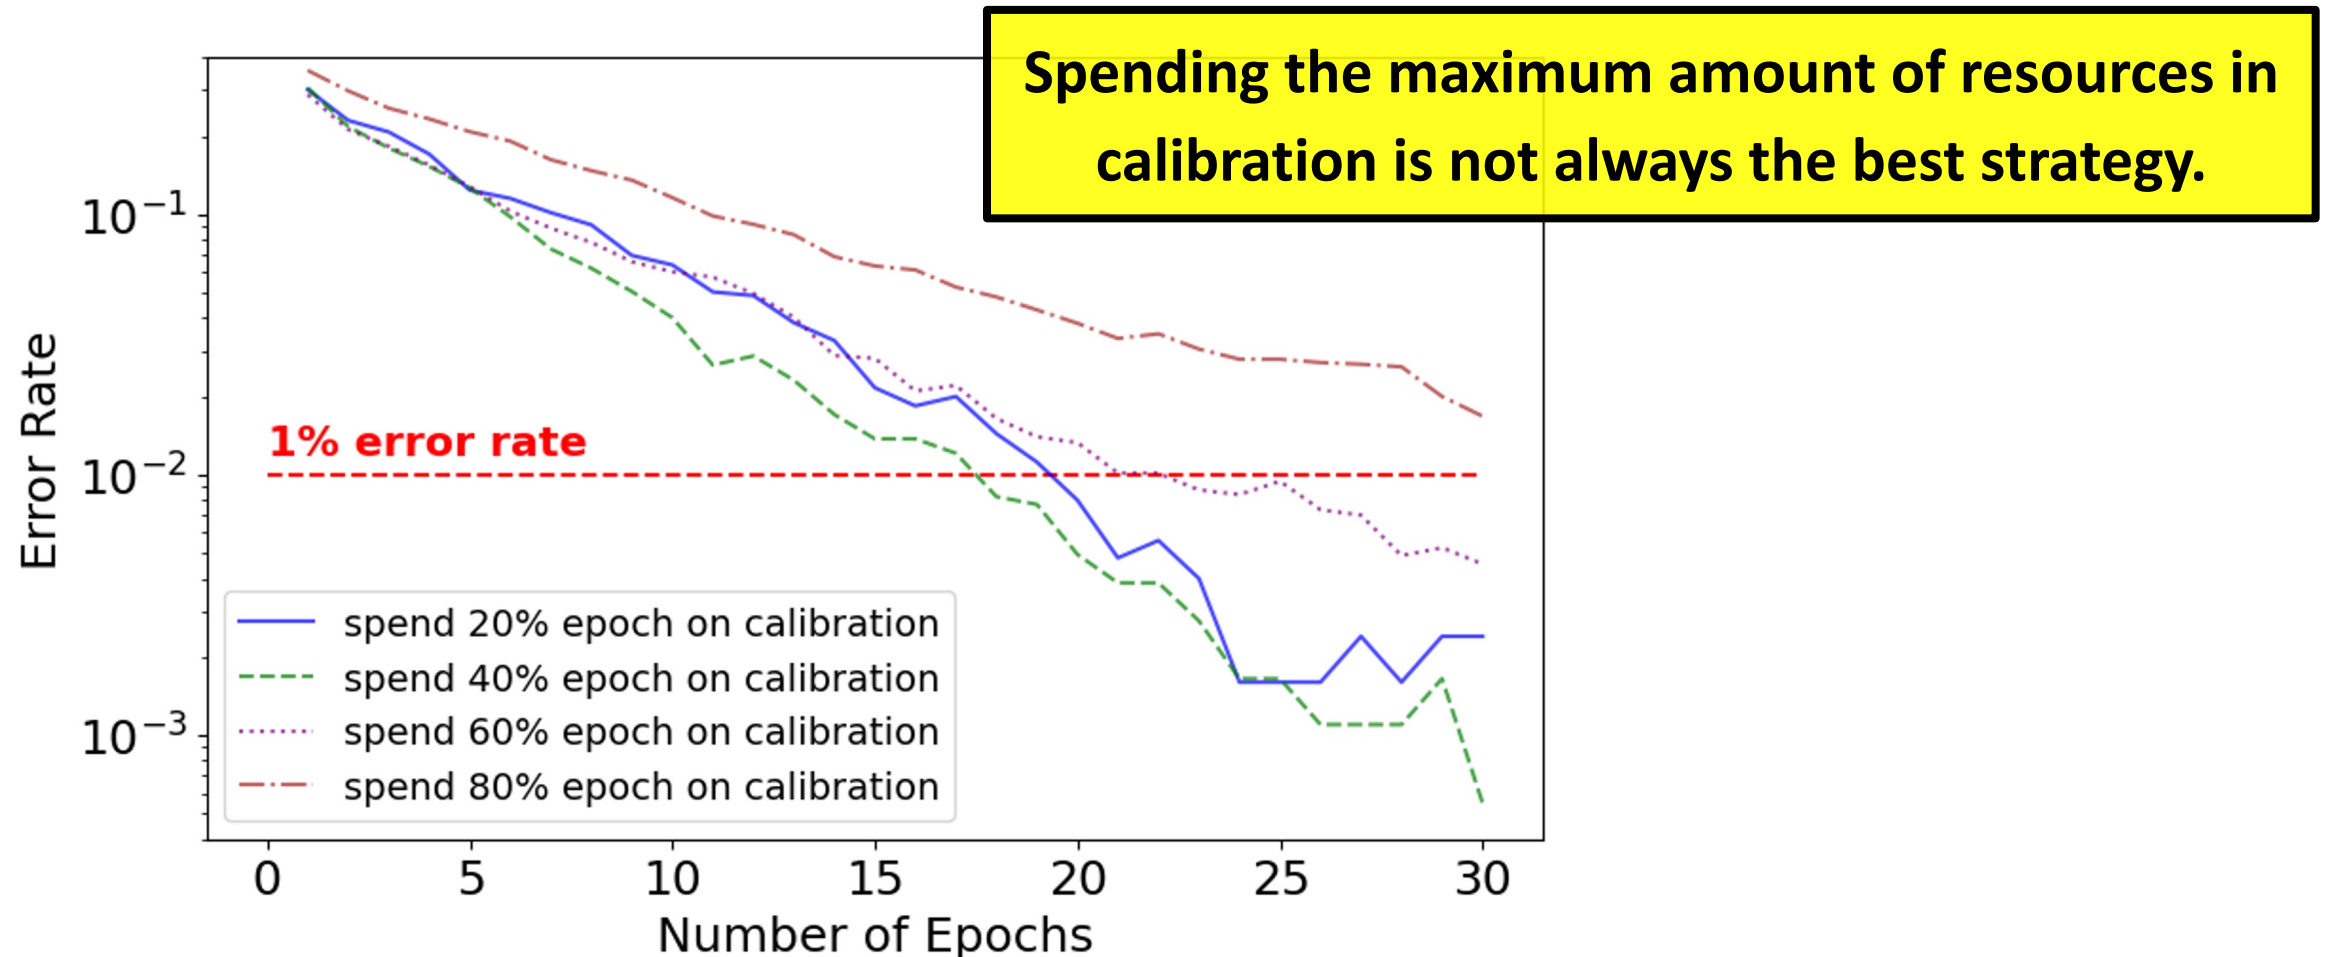


Signals across epochs when attacking the RSA square-and-multiple function.

Cross-epoch Communication


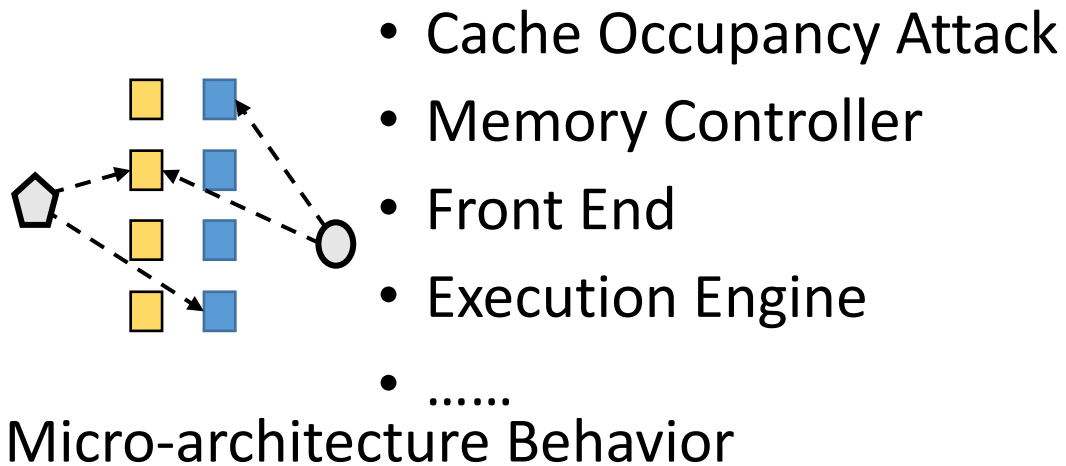


Trade-off between Calibration and Signaling



Conclusion & Long-term Impact

- Comprehensive security analysis for micro-architecture side channels should focus on end-to-end communication
- CaSA formalize the analysis of micro-architecture behavior to the analysis of random variables.



Random Variable Analysis