

# Blinded Memory

*N. Asokan, Secure Systems Group*

 <https://asokan.org/asokan/>

 @nasokan

*(Joint work with Hossam ElAtali, Lachlan J. Gunn, Hans Liljestrands)*

# This talk in a nutshell

## 1. Outsourced computing is **everywhere...**

- Machine learning models kept behind remote APIs

## 2. ...but this introduces security risks...

- Providers **don't expose models/code** to clients
- Clients **expose sensitive data** to providers
- Existing solutions like FHE/TEEs have drawbacks


## 3. ...so we propose a new solution, **Blinded Memory**

- Attestation + standard encryption + hardware-assisted taint tracking
- **Sensitive data not exposed** to output devices or covert channels

# Scenario: outsourced computation

**Goal:** run the **server's confidential code** over **client's confidential data**


- Initial target: Outsourced ML inference and/or training



**How can the client avoid revealing data to the service provider?**

- **Fully-Homomorphic Encryption:** slow due to **computational overhead**
- **Multi-Party Computation:** slow due to **network overhead**
- **Hardware-based isolation + remote attestation:** **fast**

# Hardware-assisted TEEs are pervasive



## Hardware support for

- Isolated execution: **Isolated Execution Environment**
- Protected storage: **Sealing**
- Ability to convince remote verifiers: **(Remote) Attestation**

## Trusted Execution Environments (TEEs)

Operating in parallel with “rich execution environments” (REEs)

Cryptocards



<https://www.ibm.com/security/cryptocards/>

Trusted Platform Modules



<https://www.infineon.com/tpm>

ARM TrustZone



<https://www.arm.com/products/security-on-arm/trustzone>

Intel Software Guard Extensions

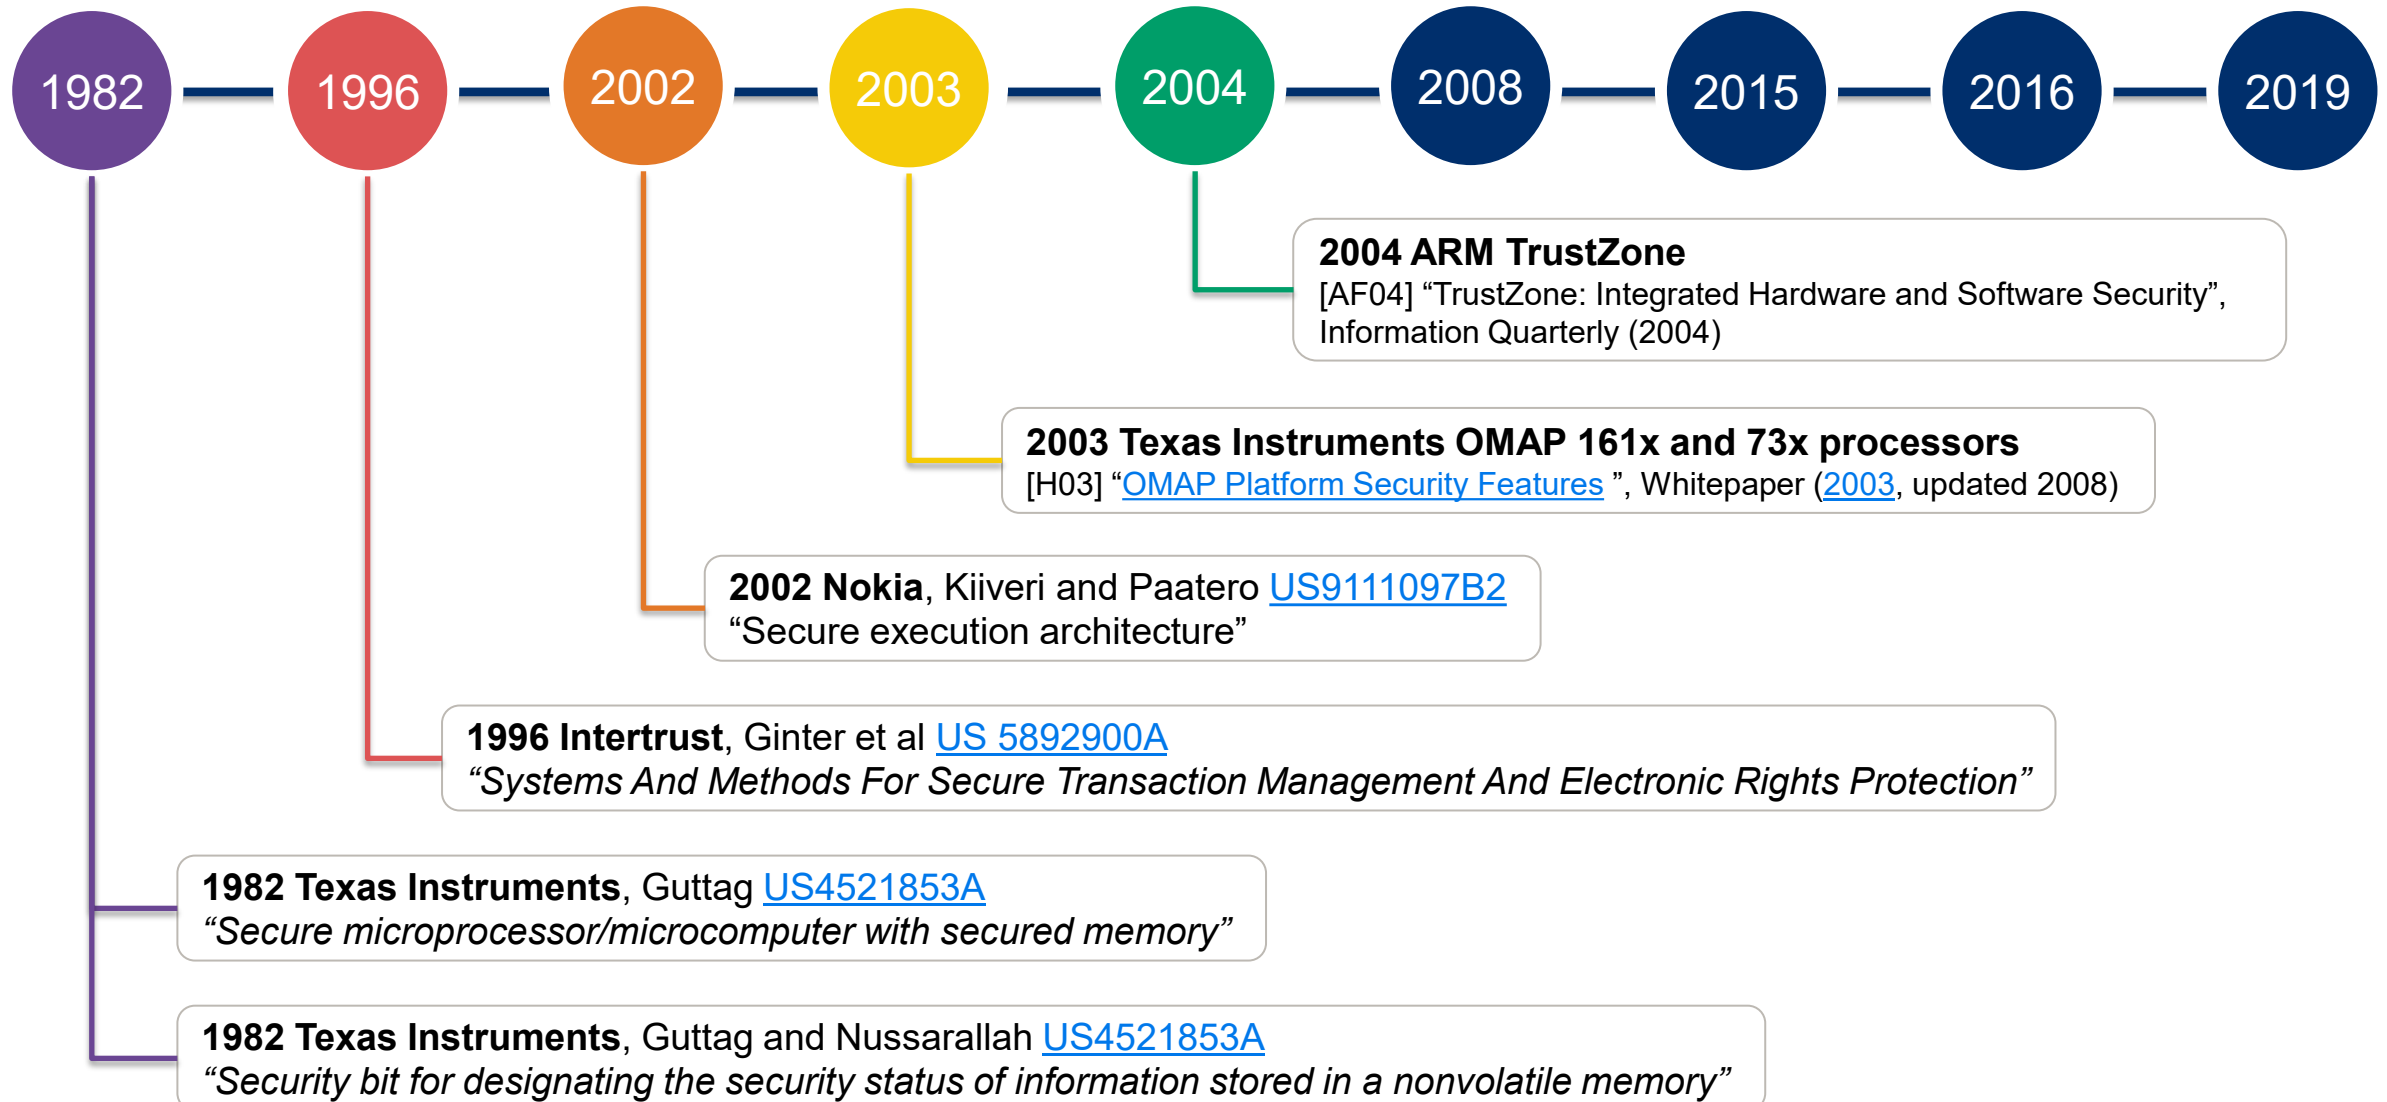


<https://software.intel.com/en-us/sgx>

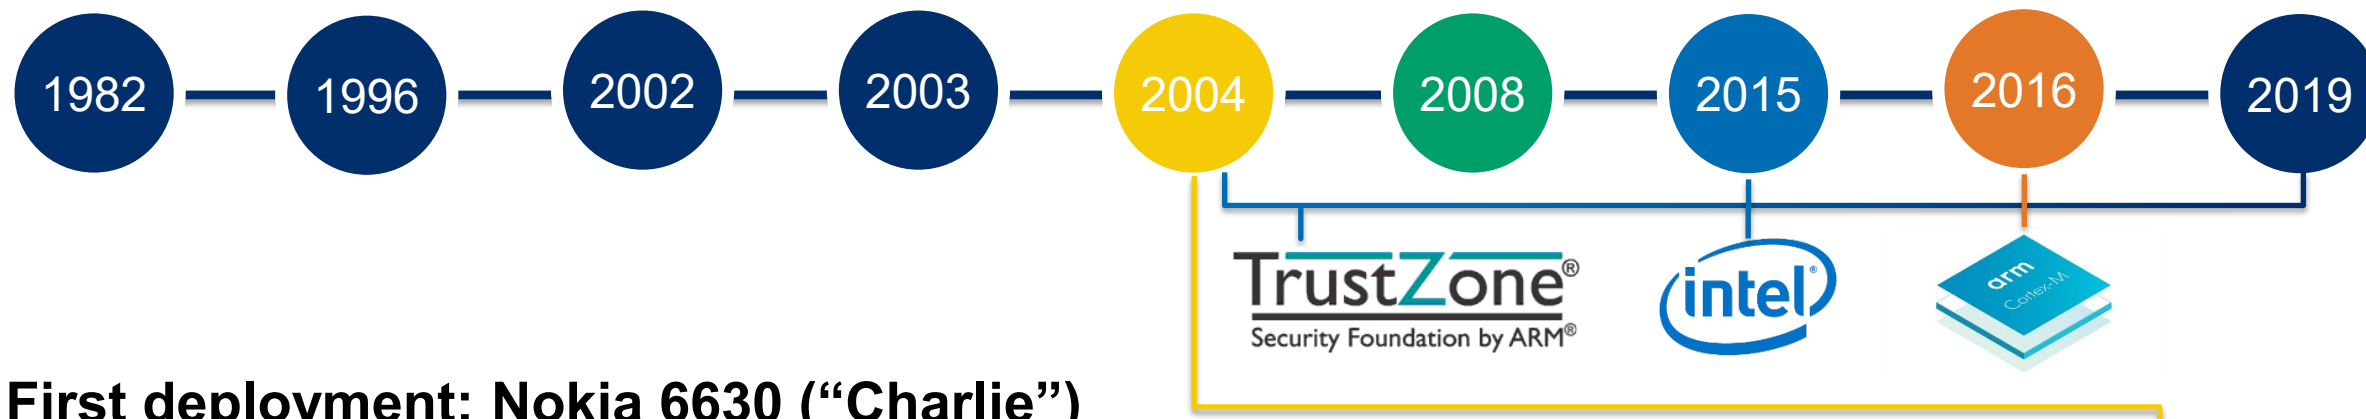
[A+14] “[Mobile Trusted Computing](#)”, Proceedings of the IEEE, 102(8) (2014)

[EKA14] “[Untapped potential of trusted execution environments](#)”, IEEE S&P Magazine, 12:04 (2014)

# TEEs as an idea date back to the 1980s



# Deployment of mobile TEEs date back to the 2000s



## First deployment: Nokia 6630 (“Charlie”)

- first 3G phone with TI OMAP 1710 processor (June 2004)

## ARM TrustZone currently widely deployed

- [TrustZone-M for Cortex-M class microcontrollers](#) (2016)

## Ca. 2008, TEE unheard of in academic circles

- first papers in FC 2008, ASIACCS 2009
  - [AE08] [A Platform for OnBoard Credentials](#), Financial Cryptography and Data Security (2008)
  - [KEAR09] [On-board credentials with open provisioning](#), ACM ASIACCS (2009)

## Intel SGX

- SkyLake (2015); wide availability of SDK “democratized” TEE research




# More on the history of TEEs

CCS 2019 keynote<sup>[1]</sup> <https://youtu.be/hHYoGn5PSI4>



2022 book <https://ssg.aalto.fi/publications/hardware-platform-security-for-mobile-devices/>



[A20] "[Hardware-assisted Trusted Execution Environments: Look Back, Look Ahead](#)", ACM CCS Keynote (2019)

[GAE+22] "[Hardware Platform Security for Mobile Devices](#)", Foundations and Trends® in Privacy and Security 3(3-4):214-394, NOW publishers (2022)

# Protection provided by TEEs comes with caveats

TEEs provide an **isolated environment** for execution of software

TEEs are **unsuitable** when server code is confidential or unverifiable

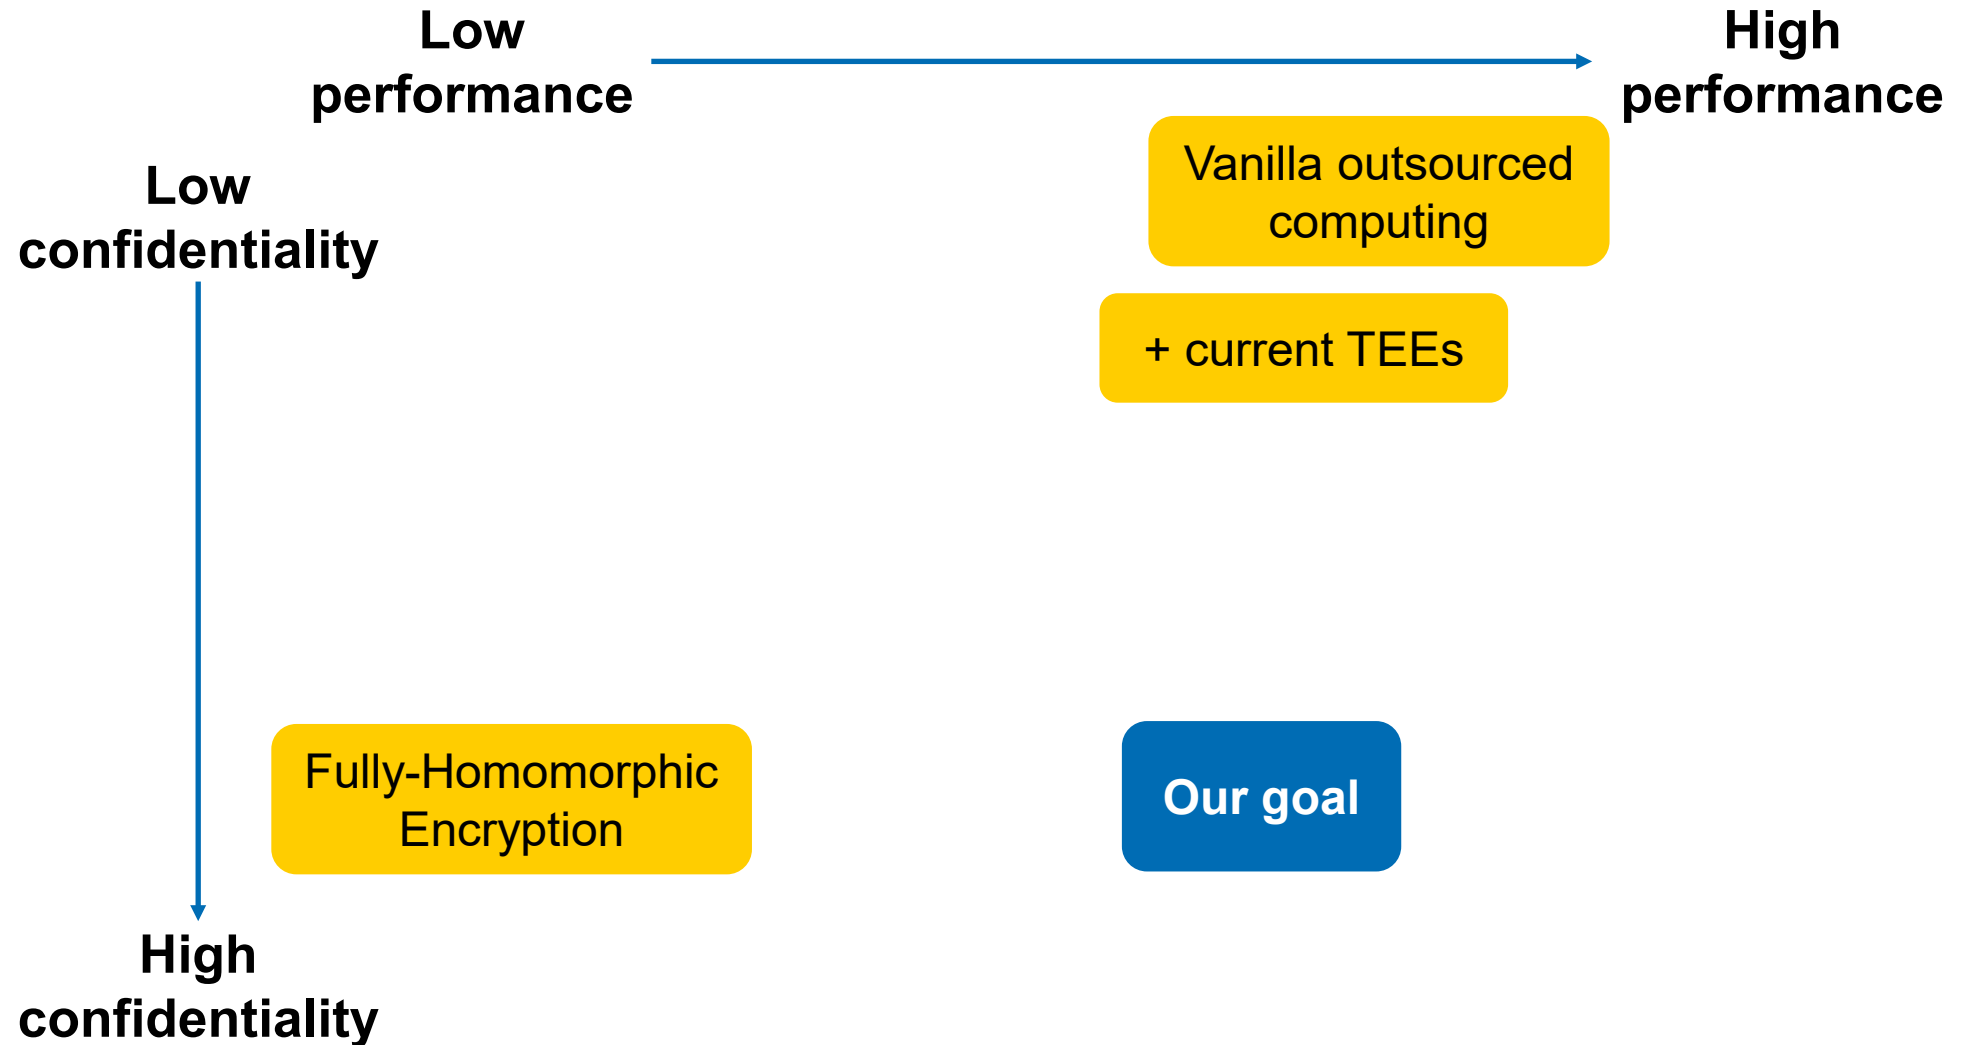
- TEEs intended for clients to run **code they trust** and can verify

Confidentiality of client data in TEEs is hampered by:

- Large TEE code base → vulnerable to **software flaws**
- Sharing resources → vulnerable to **side channels**



# Is Confidentiality vs. Performance a tradeoff?



# Prevent leakage of sensitive data via CPU extensions

“Safe” streams of instructions don’t expose **sensitive** data

## Allowed:

- Computation on sensitive data by **arbitrary, unattested, untrusted software**


## Prohibited:

- Leaking sensitive data **into any observable state**, e.g.: peripherals **outside security boundary, microarchitectural state**

Use **taint-tracking-based security policy** to limit sensitive data to **safe places**

# Combine with attestable HSM to assure clients

Remote attestation assures use of client data is **subject to security policy**




# Taint tracking policy

Registers/memory have an associated “sensitive” bit (“Blinded”

Ideal rule:

$$\text{Blinded}(\text{output}_m) \leftarrow \exists n: \text{Blinded}(\text{input}_n) \wedge \text{Depends}(\text{output}_m \text{ on } \text{input}_n)$$

**Goal:** changes in sensitive state never affect non-sensitive state  
(formally verified)



# Thinking beyond registers and memory

**Taint-propagation rule must consider many different observable outputs**

- Registers
- Memory values
- Memory access patterns
- Control flow
- Exceptions

**Not all of these outputs can be marked as sensitive**

**Data flows from sensitive values to “un-markable” outputs must yield a **fault****

# How to deal with exceptions

## Examples of **data-dependent exceptions**:

- Division by zero
- Floating-point exceptions
- ...


Instructions **must not raise an exception** based on data-dependent conditions

## Solutions:

- Unconditional faults (i.e., division by sensitive values always fails)
- **Set a sensitive error flag** and continue computation

# BliMe Architecture

1. Handshake (incl. remote attestation)
2. Shared secret key
3. Atomic data import (inputs)
  - Decrypt & blind (Blinded ← true)
4. Safe (“blinded”) computation
  - Enforced by BliMe HW extensions
5. Atomic data export (result)
  - Encrypt & unblind (Blinded ← false)



# BlMe-BOOM Implementation

On **speculative OoO RISC-V core BOOM**

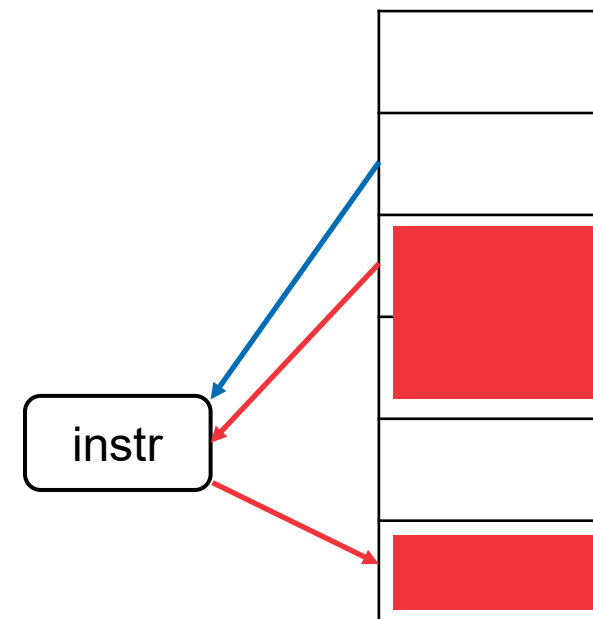
Tagged memory: each byte can be marked as **blinded**

Instructions to mark physical memory as

- **Blinded** or **non-Blinded**

Implements taint-tracking for all instructions

- Ideal rule:  $\text{Blinded}(\text{output}) \leftarrow \exists n: \text{Blinded}(\text{input}_n) \wedge \text{Depends}(\text{output on input}_n)$
- Approximate to:  $\text{Blinded}(\text{outputs}) \leftarrow \text{Blinded}(\text{input}_1) \vee \text{Blinded}(\text{input}_2) \vee \dots$



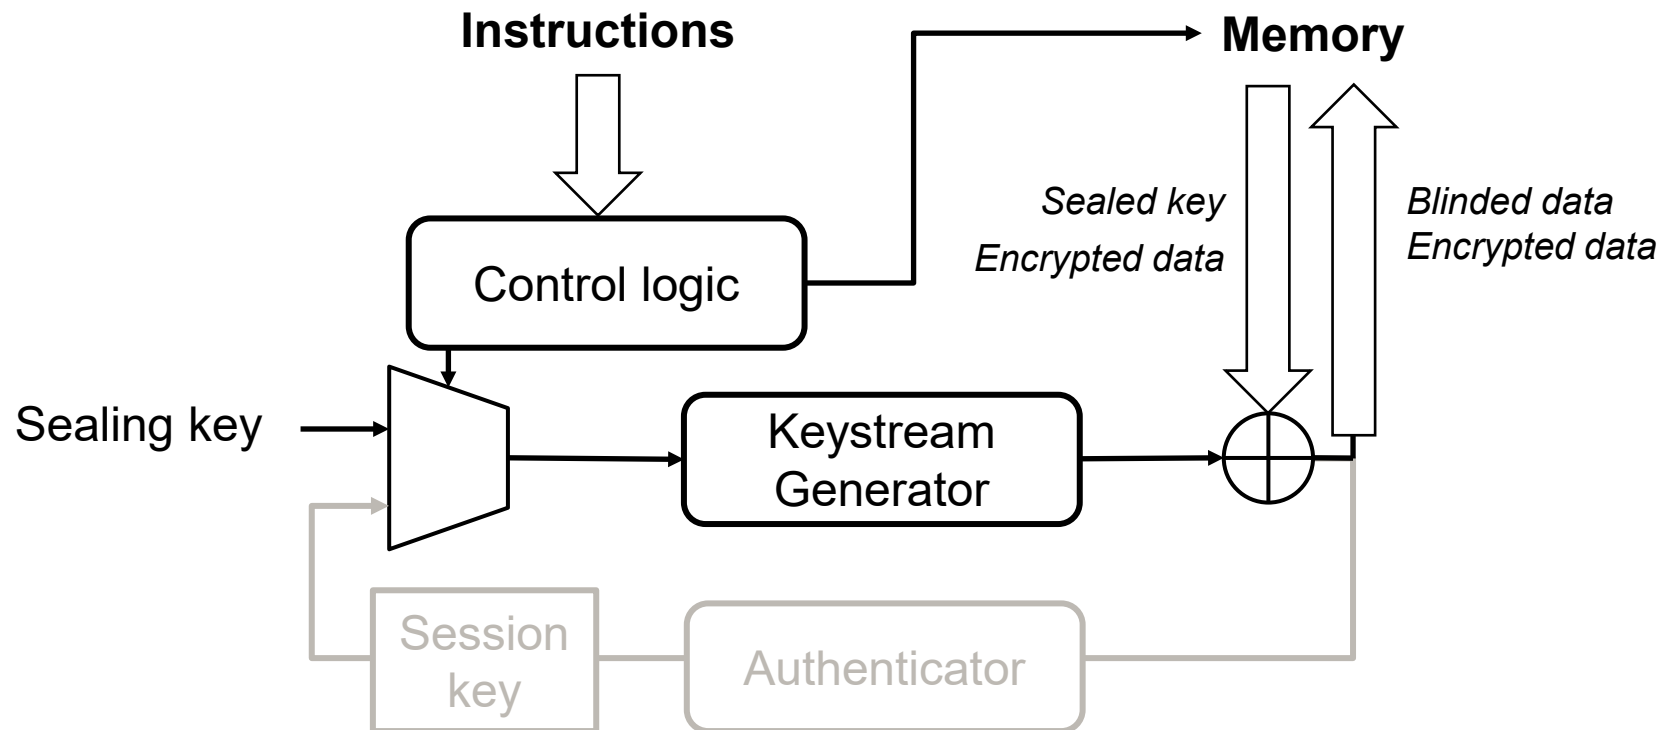
Approximation can be overridden for specific instructions



# Encryption Engine

## Encryption engine uses the RoCC accelerator interface in BOOM

- RoCC exposes custom logic as instructions



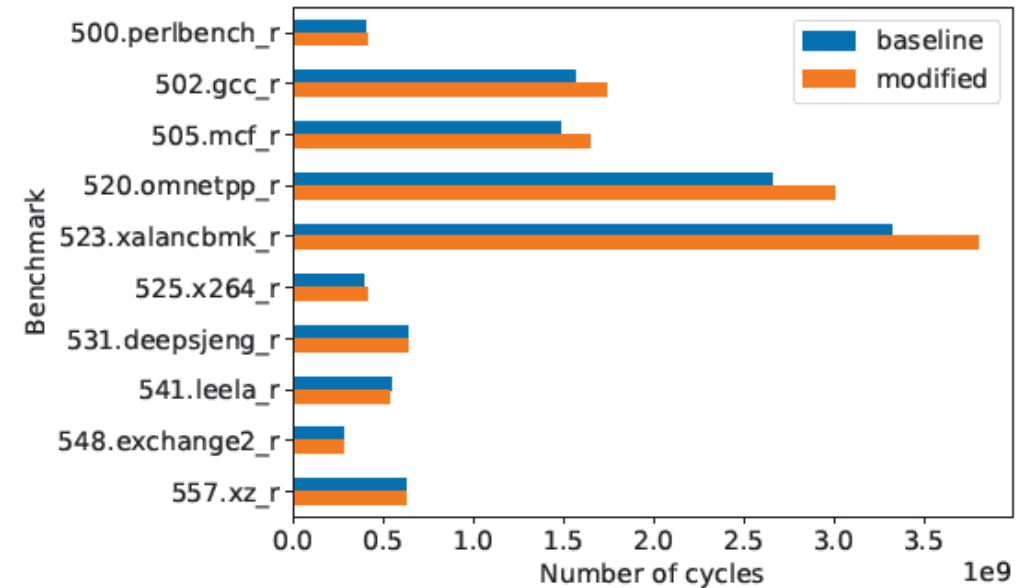
# Evaluation

## Performance: Overhead of taint-tracking logic

### FPGA

Overhead	Ibex	BOOM
LUTs & Registers	+1.3%	<+1%
Power	+2.0%	-0.2%
Max clock frequency	No reduction	No reduction

### SPEC2017 on gem5



## Compatibility: Tested with side-channel-resistant crypto library (TweetNaCl)

- Side-channel-resistant crypto runs without modifications


# Handling multiple clients simultaneously

**Problem: So far, one Blinded bit for many clients**

- Server can send sensitive data **to the wrong client**

**We need a **separate** sensitivity domain for each client**

- Prevent clients accessing each others' sensitive data
- Keys need to be swapped in and out for each client



# Handling multiple clients simultaneously

## Solution 1: Blime-BOOM-1 + Isolation by honest-but-curious server OS

- OS keeps track of sensitivity domains
- Requires only **single Blinded bit** from HW: **low memory overhead**
- Rely on remote **attestation of the entire OS** to convince client

## Solution 2: BliMe-BOOM-N -- Hardware support for multiple clients

- Hardware keeps track of sensitivity domains: **multibit Blindedness indicator**
- Secure **despite malicious OS**
- Requires **extra memory/logic** to keep track of domain identifier for each page

# BliMe-BOOM-N Implementation (eval in progress)

**BOOM RTL**

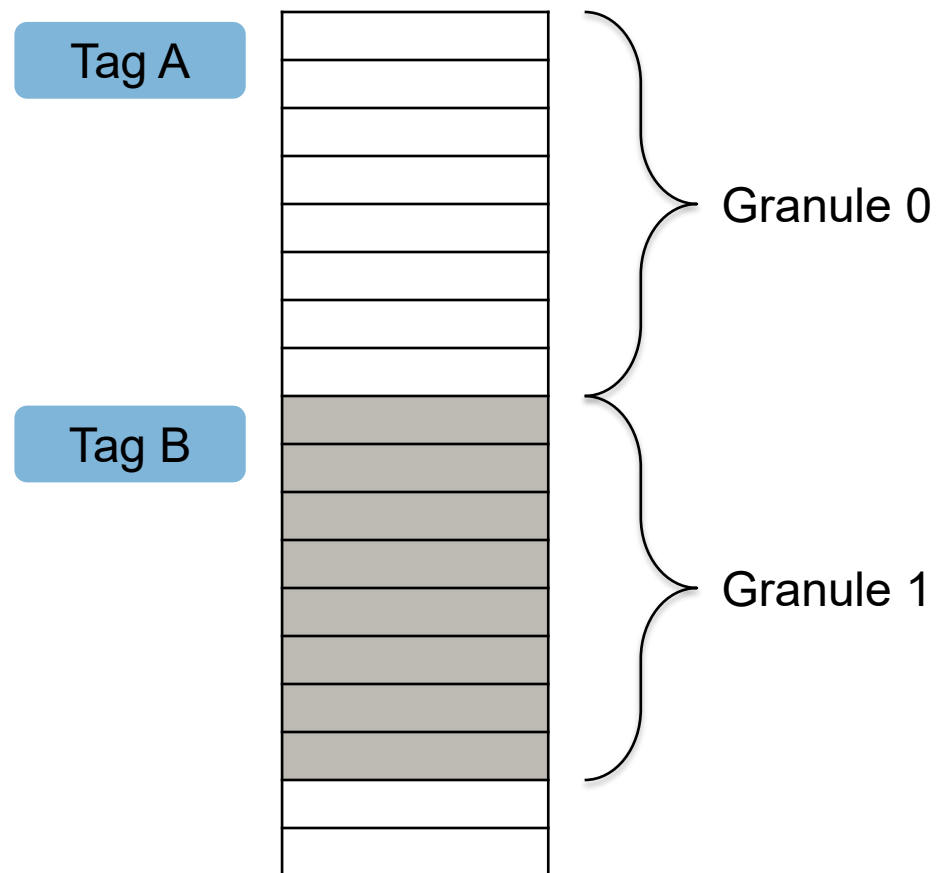
Data tagged with **client-specific tag**

**1 tag per granule**

**Tag size = 8 bits, granule size = 8 bytes**

**Future work:**

- parameterize tag size and granule size



# Security: Formal verification in F\*

## Flag-based error handling

Challenge: common error-handling patterns have some “leakage”

E.g., side-channel-resistant authenticated encryption, faults on BlIME

- Decryption control flow depends on authentication check result
- Authentication check result depends on (sensitive) key

Solution: modify APIs to produce [meaningless?] output + an error flag

- Computation continues, ignoring errors
- Error flag returned to the client, who can ignore the data if an error occurred

**Goal: changes in blinded state never affect non-blinded state**

```
(*****  
 * Equivalence-based safety.  
 *  
 * We define safety in this case to be that the system is safe if executing on  
 * equivalent (and so indistinguishable) states always results in equivalent  
 * output states.  
 *****)  
let equivalent_inputs_yield_equivalent_states (exec:execution_unit) (pre1 pre2 : systemState) =  
    equiv_system pre1 pre2 ⇒ equiv_system (step exec pre1) (step exec pre2)  
  
let is_safe (exec:execution_unit) =  
    ∀ (pre1 pre2 : systemState). equivalent_inputs_yield_equivalent_states exec pre1 pre2
```

# Flag-based error handling

**Challenge: common error-handling patterns have some “leakage”**

**E.g., side-channel-resistant authenticated encryption, faults on BliMe**

- Decryption control flow **depends on authentication check result**
- Authentication check result **depends on (sensitive) key**

**Solution: modify APIs to produce [meaningless?] output + an error flag**

- Computation continues, **ignoring errors**
- Error flag **returned to the client**, who can **ignore the data** if an error occurred

# Generating compliant code with LLVM

## Summary

BliMe provides FHE-style security, but **efficiently**

Server can **safely run untrusted code** on sensitive data

Incorporated into **speculative OoO RISC-V core BOOM**



[EGH22] "BliMe: Verifiably Secure Outsourced Computation with Hardware-Enforced Taint Tracking", arXiv:2204.09649 [cs.CR] (2022)

More on our research at <https://crisp.uwaterloo.ca/research/SSG/>  
<https://ssg.aalto.fi/research/projects/platsec/>

## Problem: software **might not run as-is**

- BliMe hardware extensions will abort non-compliant code

## Creating **compliant code by hand is error prone**

- High-level verification often insufficient
- Challenge exacerbated due to **obtuse compiler behavior**
- **Usability challenge**, not **security**

## Solution: **taint-tracking compiler**<sup>[B+21]</sup>

- Ensure **policy-compliance** during code compilation
- Transform **non-compliant code** to **compliant equivalent**

[B+21] "[Constantine: Automatic Side-Channel Resistance Using Efficient Control and Data Flow Linearization](#)", ACM CCS (2021)<sup>30</sup>



# Compiler taint tracking

**Problem: Constantine's profiling under-approximates taint**

- Can cause **incorrect faults** with BliMe - unacceptable

**Solution: Use static analysis to propagate taint**

- Trade-off: over-approximation

**Use SVF<sup>[S+16]</sup> as a starting point**

**SVF provides static value-flow graph**

- Shows value dependencies within program

# BliMe-X

## HW accelerators common in outsourced computation

- Customized to application

## ML accelerators useful for all ML workloads

## Goal: Adapt BliMe to ML accelerator framework

## Prominent open-source frameworks:

- NVDLA (by NVIDIA)
- Gemmini (by BOOM team)




# Summary

BliMe provides FHE-style security, but **efficiently**

Server can **safely run untrusted code** on **sensitive data**

Incorporated into **speculative OoO RISC-V core BOOM**



[EGHA22] "[BliMe: Verifiably Secure Outsourced Computation with Hardware-Enforced Taint Tracking](https://arxiv.org/abs/2204.09649)", arXiv:2204.09649 [cs.CR] (2022)

More on our research at <https://crisp.uwaterloo.ca/research/SSG/>  
<https://ssg.aalto.fi/research/projects/platsec/>