



Aalto University

# Common-sense applications of hardware-based TEEs


*N. Asokan*

 <http://asokan.org/asokan/>

 [@nasokan](https://twitter.com/nasokan)

*Acknowledgements: Thomas Nyman, Lachlan Gunn*

# Hardware-security mechanisms 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




Trusted Platform Modules



ARM TrustZone



Intel Software Guard Extensions



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

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

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

<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)

# Concerns with TEEs: flaws

## TPM Reset Attack

50,012 views



Evan Sparks

Published on Jun 18, 2007

A demonstration of a vulnerability in the TCG architecture v running TPM without restarting the platform.

<http://www.cs.dartmouth.edu/~pkilab/sparks/> (2007)

## CLKSCREW: Exposing the Perils of Security-Oblivious Energy Management

Authors:

Adrian Tang, Simha Sethumadhavan, and Salvatore Stolfo, *Columbia University*

*Distinguished Paper Award Winner!*

<https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/tang> (2017)

## Foreshadow (security vulnerability)

From Wikipedia, the free encyclopedia

*This article is about the security vulnerability. For other uses, see Foreshadow (disambiguation).*

**Foreshadow** is a vulnerability that affects modern microprocessors that was first discovered by two independent teams of researchers in January 2018, but was first disclosed to the public on 14 August 2018.<sup>[1][2][3][4][5][6][7][8][9][10][11][12][13][14][15][16]</sup> The vulnerability is a speculative execution attack on Intel processors that may result in the loss of sensitive information stored in personal computers, or third party clouds.<sup>[1]</sup> There are two versions: the first version (original/Foreshadow) (CVE-2018-3615<sup>[#]</sup>) targets data from SGX enclaves; and the second version (next-generation/Foreshadow-NG<sup>[#]</sup>) (CVE-2018-3620<sup>[#]</sup> and CVE-2018-3646<sup>[#]</sup>) targets Virtual Machines (VMs), hypervisors (VMM), operating system (OS) kernel memory, and System Management Mode (SMM) memory.<sup>[1]</sup> Intel considers the entire class of speculative execution side channel vulnerabilities as "L1 Terminal Fault" (L1TF).<sup>[1]</sup> A listing of affected Intel hardware has been posted.<sup>[10][11]</sup>

Foreshadow is similar to the Spectre security vulnerabilities discovered earlier to affect Intel and AMD chips, and the Meltdown vulnerability that also affected Intel.<sup>[6]</sup> However, AMD products, according to AMD, are not affected by the Foreshadow security flaws.<sup>[6]</sup> According to one expert, "[Foreshadow] lets malicious software break into secure areas that even the Spectre and Meltdown flaws couldn't crack".<sup>[15]</sup> Nonetheless, one of the variants of Foreshadow goes beyond Intel chips with SGX technology, and affects "all [Intel] Core processors built over the last seven years".<sup>[2]</sup>

Foreshadow may be very difficult to exploit.<sup>[2][6]</sup> and there seems to be no evidence to date (15 August 2018) of any serious hacking involving the Foreshadow vulnerabilities.<sup>[2][6]</sup> Nevertheless, applying software patches may help alleviate some concern(s), although the balance between security and performance may be a worthy consideration.<sup>[5]</sup> Companies performing cloud computing may see a significant decrease in their overall computing power; individuals, however, may not likely see any performance impact, according to researchers.<sup>[9]</sup> The real fix, according to Intel, is by replacing today's processors.<sup>[5]</sup> Intel further states, "These changes begin with our next-generation Intel Xeon Scalable processors (code-



[https://en.wikipedia.org/wiki/Foreshadow\\_\(security\\_vulnerability\)](https://en.wikipedia.org/wiki/Foreshadow_(security_vulnerability)) (2018)

# Concerns with TEEs: suspicions of motives

## Software

### MS Palladium protects IT vendors, not you – paper

Anderson gives us the FAQs

By John Lettice 28 Jun 2002 at 10:27

SHARE ▼

[https://www.theregister.co.uk/2002/06/28/ms\\_palladium\\_protects\\_it\\_vendors/](https://www.theregister.co.uk/2002/06/28/ms_palladium_protects_it_vendors/) (2002)

### Trusting Intel – Next Generation of Backdooring?

We have seen that SGX offers a number of attractive functionality that could potentially make our digital systems more secure and 3<sup>rd</sup> party servers more trusted. But does it really?

The obvious question, especially in the light of recent revelations about NSA backdooring everything and the kitchen sink, is whether Intel will have backdoors allowing “privileged entities” to bypass SGX protections?

<http://theinvisiblethings.blogspot.fi/2013/09/thoughts-on-intels-upcoming-software.html> (2013)

**Problem: Third-party uncertainty about your software environment is normally a feature, not a bug**

<https://www.eff.org/wp/trusted-computing-promise-and-risk> (2003)

# Possible motivations for widespread deployment

**Vendor lock-in**

**Restriction of digital rights**

...

**Regulatory requirements**

**Protection of end-user data**

...


# Example: regulatory compliance

The IMEI shall not be changed after the ME's final production process. It shall resist tampering, i.e. manipulation and change, by any means (e.g. physical, electrical and software).



NOTE: This requirement is valid for new GSM Phase 2 and Release 96, 97, 98 and 99 MEs type approved after 1<sup>st</sup> June 2002.

3GPP TS 42.009, 2001

**Secure storage of RF  
configuration parameters**



**Early TEEs for mobile phones**  
(ca. 2001)




**TrustZone**<sup>®</sup>  
Security Foundation by ARM<sup>®</sup>

# Mobile TEEs: Motivation

## Business requirements:

- mobile payment
- subsidy lock
- custom silicon consolidation



## Regulatory requirements:

- tamper-resistant IMEIs
- secure storage for RF



## Supply-chain constraints:

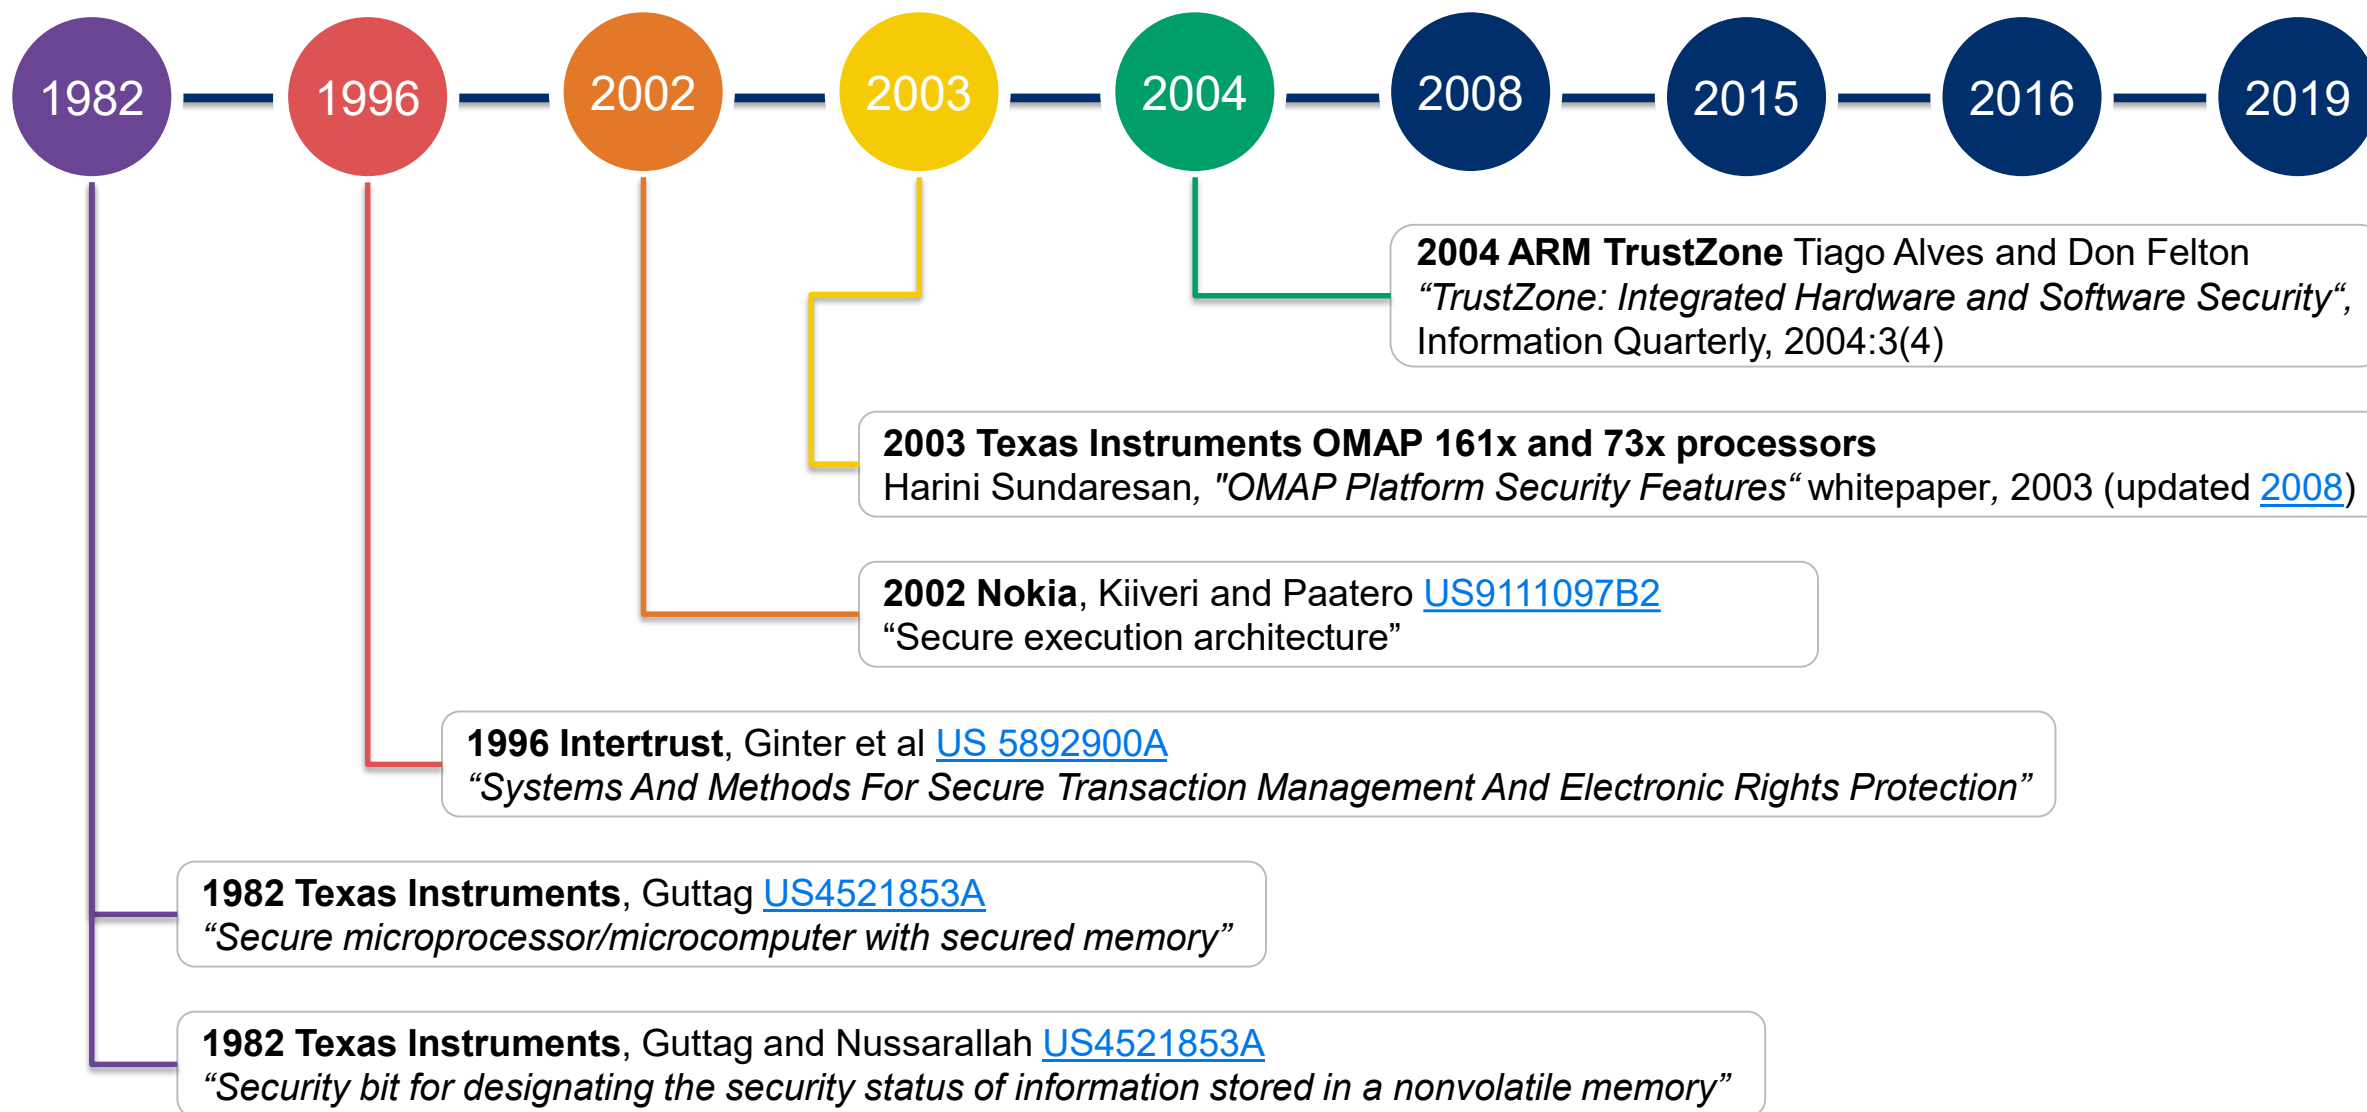
**Cost of discrete security chip  
too high on bill of materials!**



➡ New approach: “*processor secure environments*”

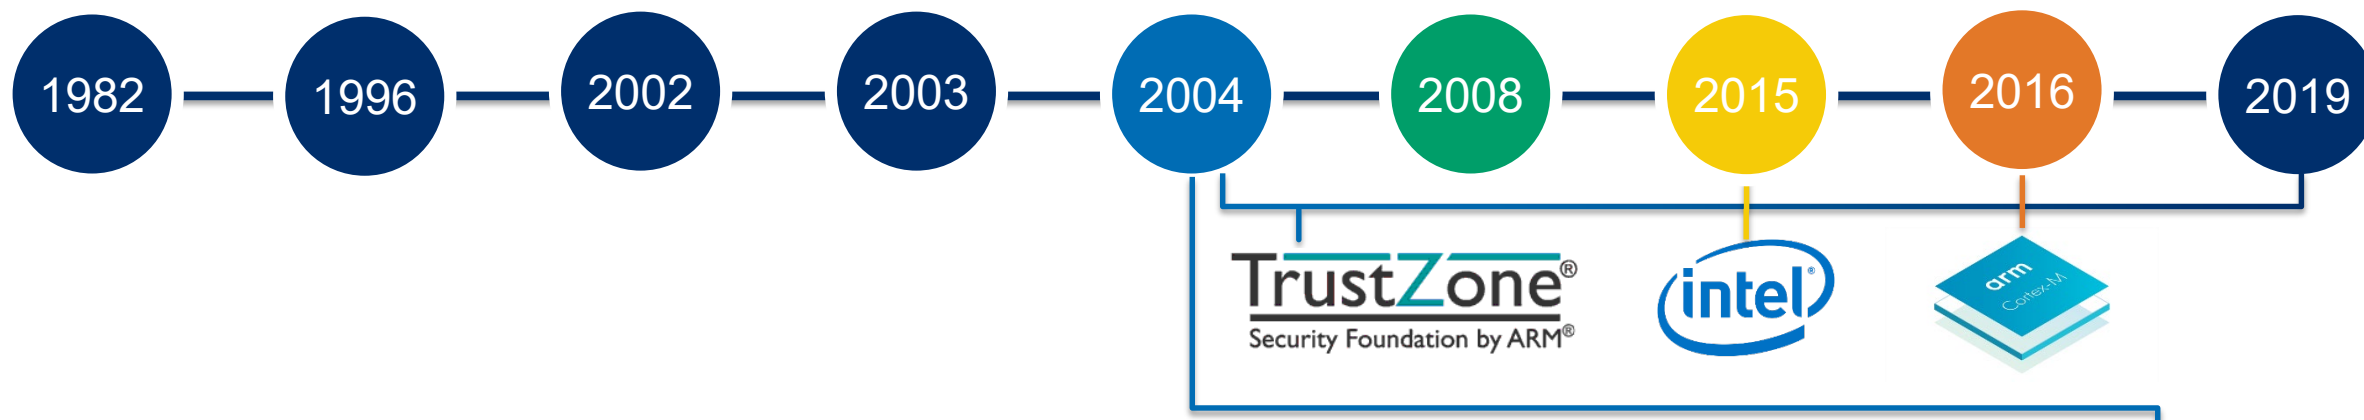
**Generic low-cost** enabler emerged as **skunkworks project** within Nokia  
(rather than point solutions for particular use cases)

# Mobile TEEs: Development





# Mobile TEEs: Deployment



- **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 academic circles**
  - first paper in FC 2008, ASIACCS 2009
- **Intel SGX**
  - SkyLake microarchitecture (2015)
  - wide availability of SDK “democratized” TEE research



# Should we build systems that rely on TEEs?

Concerns with applicability of hardware-supported TEEs remain

But compelling common-sense applications exist

**practical**; protect **end-users**; address **everyday** needs

- Private membership test for **malware scanning**, private contact discovery,..  
[TLPEPA17] Circle Game, ACM ASIACCS <https://arxiv.org/abs/1606.01655>
- Protection of **password-based web authentication**  
[KKPMA18] SafeKeeper, WWW (WebConf) <https://ssg.aalto.fi/research/projects/passwords/>
- Secure **accounting for function-as-a-service (FaaS) settings**  
[AAKPS18], S-FaaS, in submission, <https://export.arxiv.org/abs/1810.06080>
- *Blockchains and cryptocurrencies*  
[LLKA19] FastBFT, IEEE TC <https://doi.org/10.1109/TC.2018.2860009>, [GLVA19] SACZyzyyva, SRDS, <http://arxiv.org/abs/1905.10255>
- ...

# Can blockchains be made better using hardware-assisted security?

*Lachlan J. Gunn, N. Asokan*


# Proof of Work + “longest chain” rule

Bitcoin, Ethereum, etc. all use Proof of Work to agree on the next block:

Miners decide which transactions include in their proposal for the next block

Proof of Work: use computation power to solve a puzzle; winner proposes next block

- Chance of success proportional to amount of computation (work) performed
- Fair: any miner expending the same amount of work has the same chance of winning



- Everyone follows the longest valid chain (chain with largest CPU power wins eventually)

# What's wrong with Bitcoin, anyway?

The luxury of not trusting anyone does not come for free:

All transactions need to be **online**

**Slow:** long confirmation time, low throughput

**Wasteful** (energy expended on puzzle solving)

**Probabilistic** finality

Extremely **scalable**

Annual Power  
Consumption



Data: Digiconomist, CIA World Factbook

# Outline

## Can hardware-assisted security improve blockchains?

### Example approaches

- Changing the “business process”
- Replacing consensus (“longest chain” rule)
- ...

### What challenges arise?

# Changing the process

# Proof of Elapsed Time

## Proof of Work:

First miner to **solve puzzle** wins (gets to propose next block)

**Work  $\sim$  Exp (difficulty)**

*Proposals can be made at a rate proportional to computational power*

## Proof of Elapsed Time:

TEE issues **attestation** after waiting (idly) for a while; First miner to get the attestation wins

**Idle wait time  $\sim$  Exp (difficulty)**

*Proposals can be made at a rate proportional to the number of **idle** CPUs*



# Replacing Consensus

# Byzantine Consensus

## Goals of classical Consensus schemes:

- Liveness: all (honest) nodes produce **output**
- Safety: all (honest) nodes output **same value**
- Finality: output values are **definitive**

## Adversary model:

- Adversary can compromise some nodes
- Goals hold **despite  $f$  compromised nodes**


## Limits:

- **No protocol** can tolerate more than a third of nodes being compromised

Slow

Probabilistic

Wasteful



# PBFT

Fast

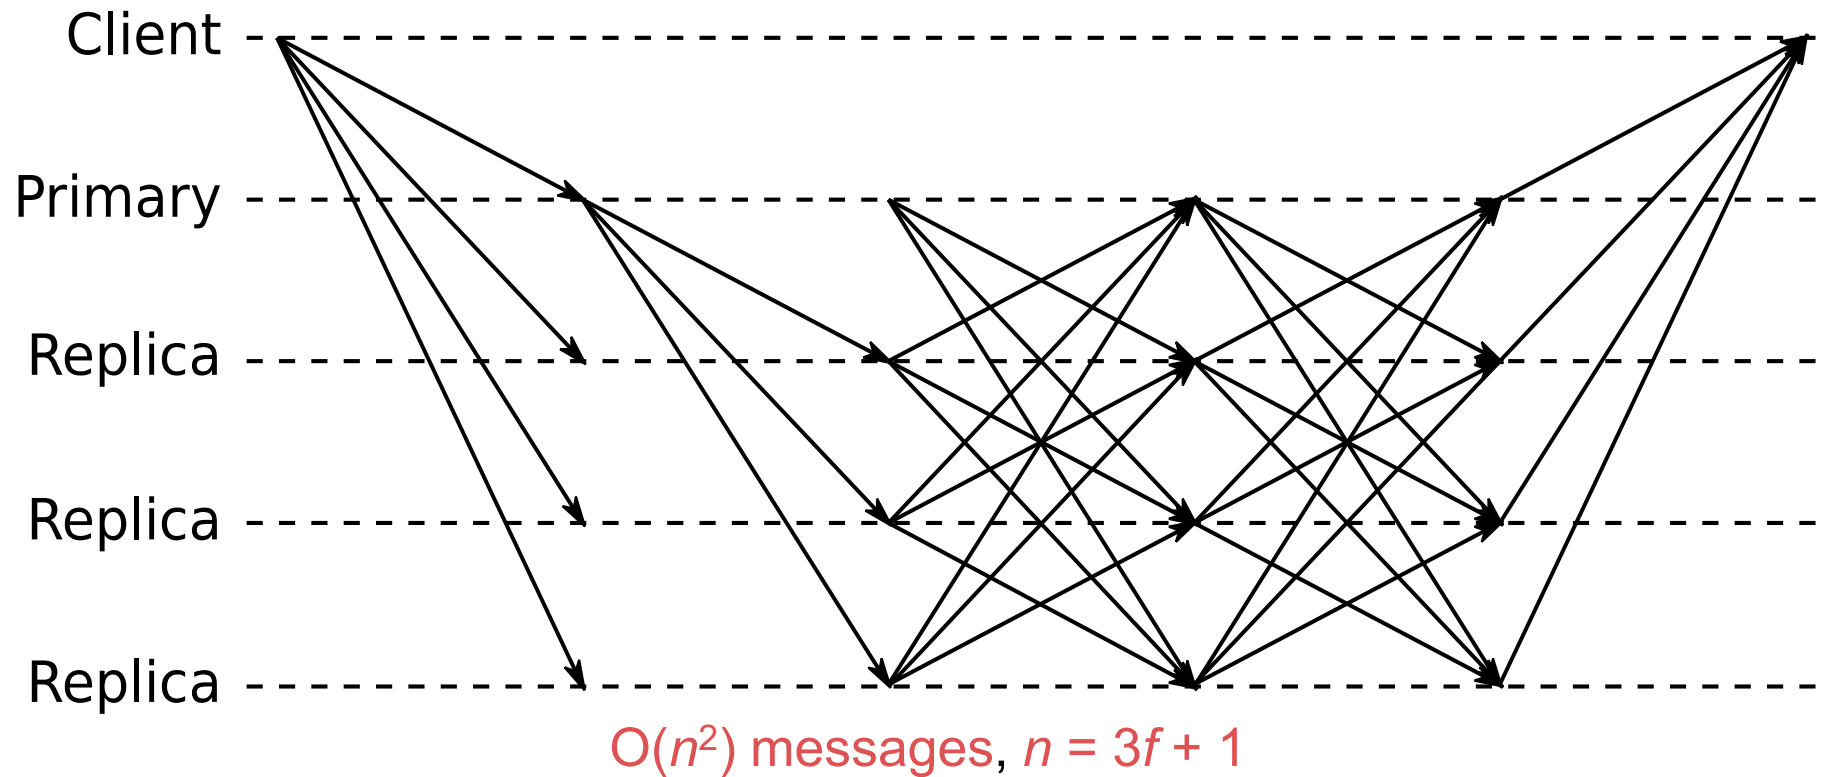
Deterministic

Efficient

Scalable

The first practical protocol for Byzantine fault tolerance

Less scalable than Proof of Work.



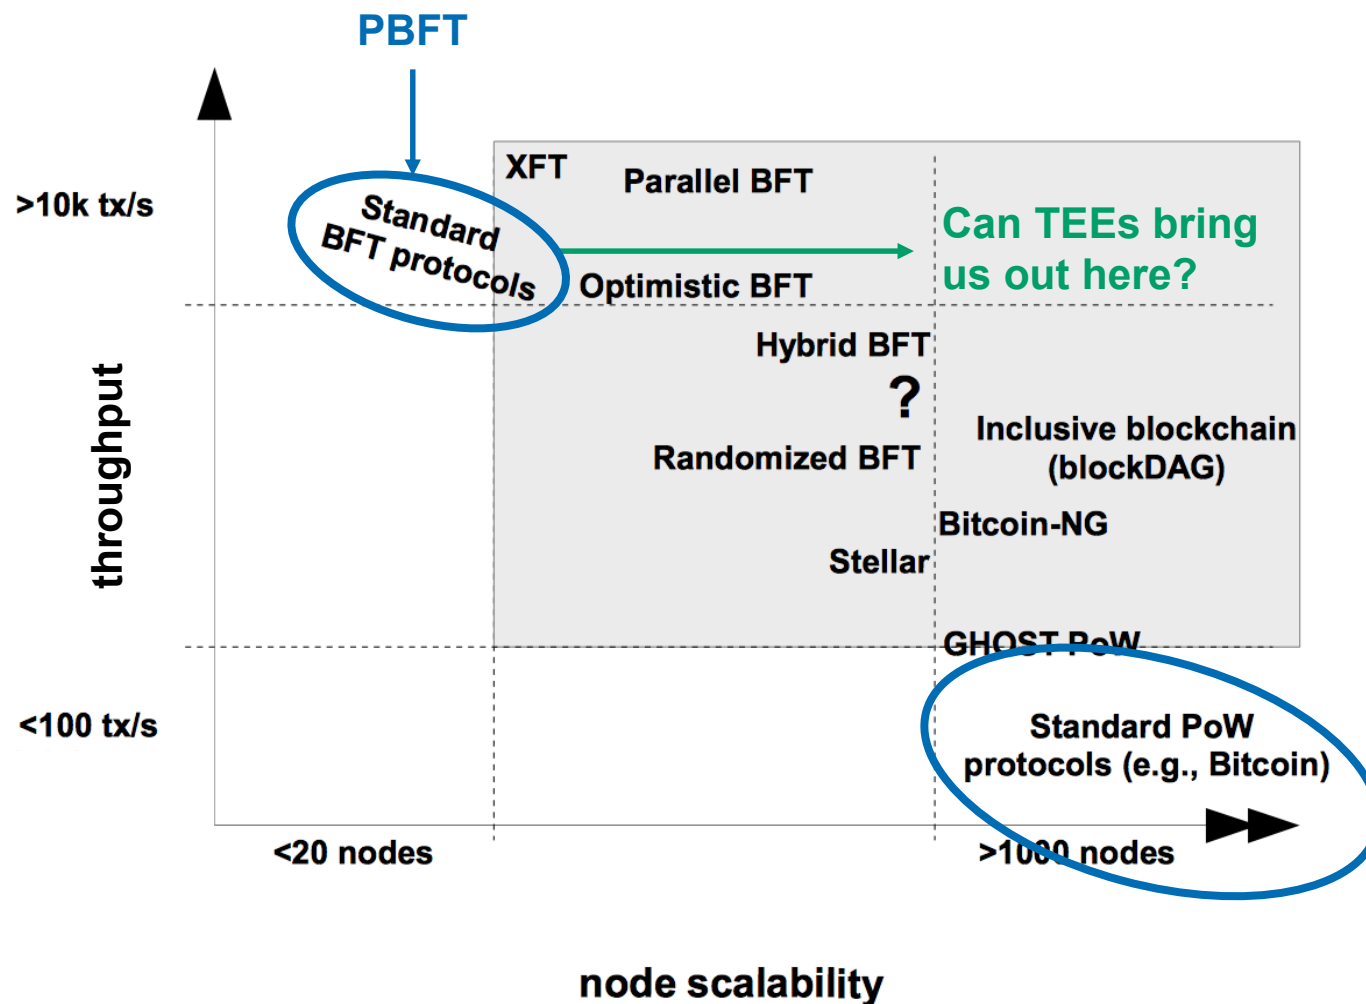
# The landscape of consensus mechanisms

Fast

Deterministic

Efficient

Scalable



Adapted from Marko Vukolić, "[The quest for scalable blockchain fabric: Proof-of-work vs. BFT replication](#)"  
International Workshop on Open Problems in Network Security. Springer International Publishing, 2015


# How can TEEs help design scalable consensus?

**Problem:** Compromised nodes can **equivocate**

**Solution:** Use **attestation** to **prevent equivocation!**


- Tolerate faults in  $\frac{1}{2}$  of the nodes

Applicability **limited to permissioned settings**




# MinBFT

PBFT



## Hardware-based monotonic counters

→ increase fault-tolerance




# FastBFT

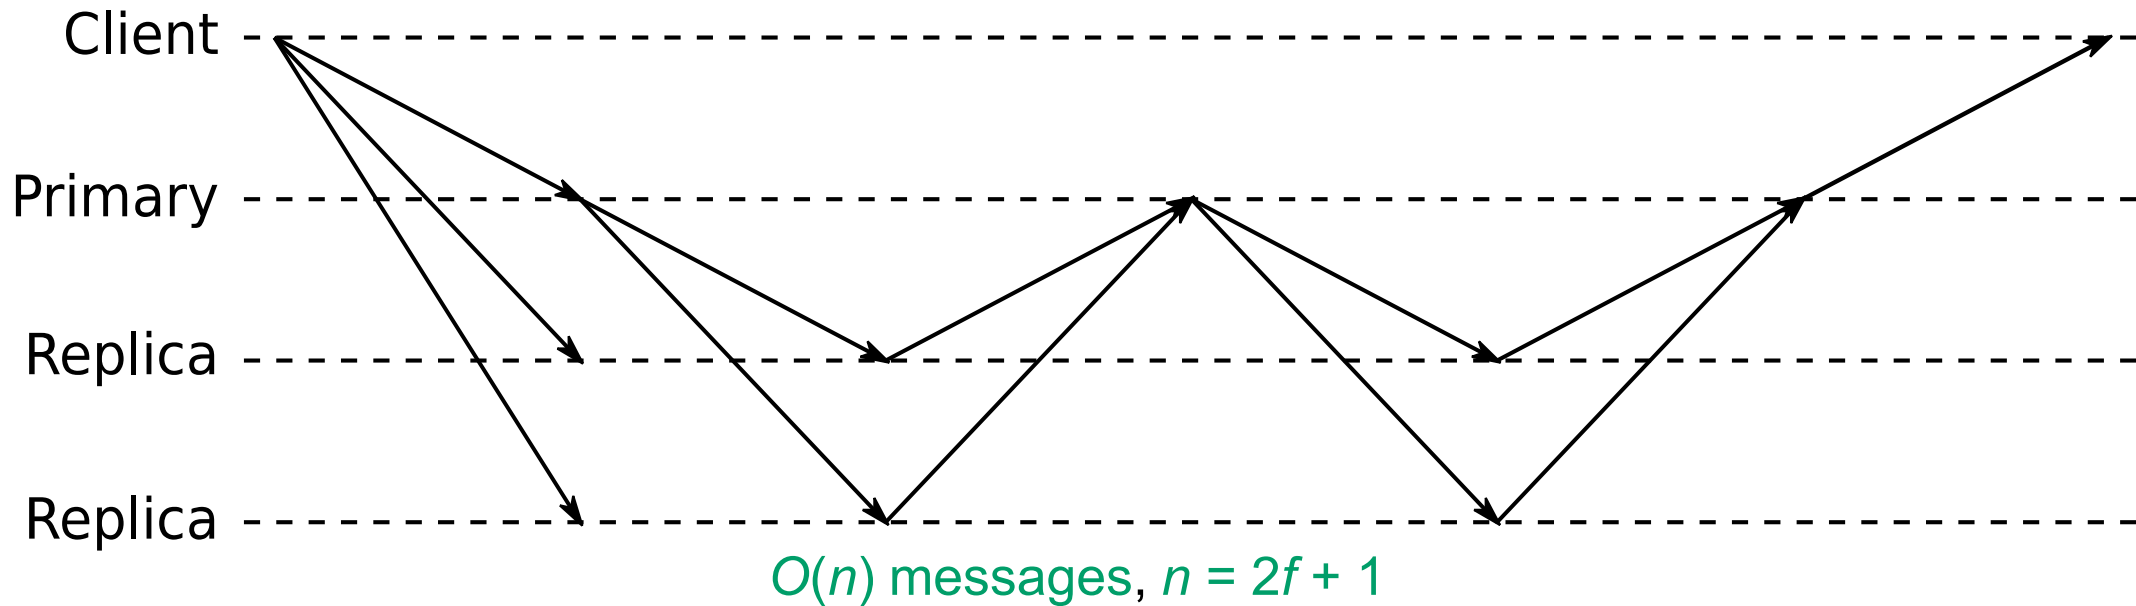

TEE-protected secret sharing, message aggregation

→ increase throughput

PBFT



MinBFT



# Challenges




# Challenges in relying on hardware-assistance

TEE unavailable


## TEE Availability:

- TEEs will not be universally available:
  - Gradual rollout
  - Obsolescence
  - Revocation



## TEE Compromise:

- Compromising some TEEs should not completely break the system



# Example: Dealing with TEE availability in consensus

**Question:** Can we improve consensus protocols by adding **only a few** TEEs?

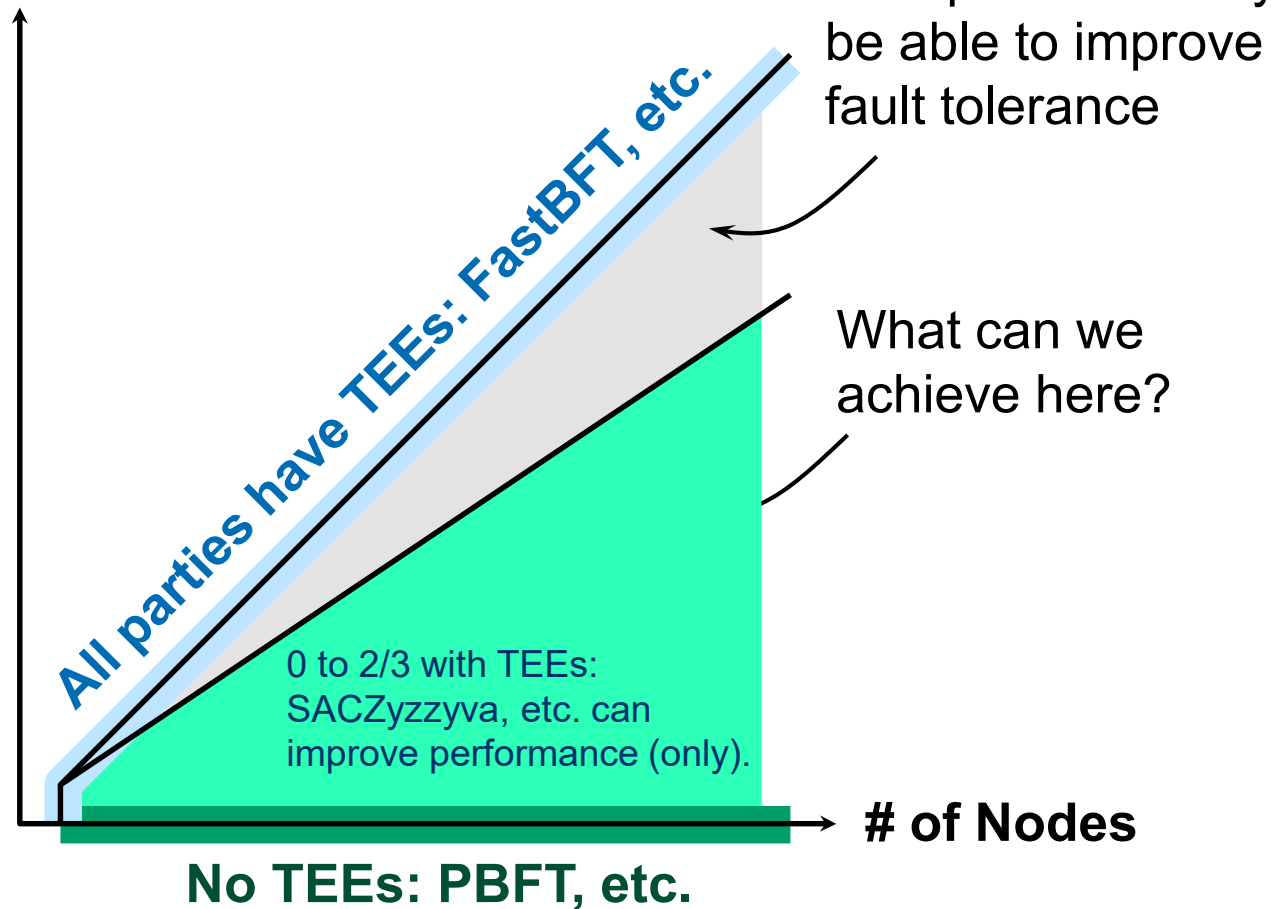
**Answer\*:**

- can **increase throughput** if  $\#TEEs > 1$
- but **fault tolerance cannot be increased** if  $(\#TEEs / \#Nodes) \leq 2/3$

**Open question:** (How) can we optimally increase fault tolerance when

$$2/3 < (\#TEEs / \#Nodes) < 1$$

# of TEEs



\* [GLVA19] SACZyzyva, SRDS, <http://arxiv.org/abs/1905.10255>

# Example: Dealing with TEE compromise in PoET


**Problem:** A **compromised TEE** can **win every block**

**Statistical solution:** refuse blocks from machines that have won too many times

- Before: compromised TEEs give attacker **unlimited power**
- After: attacker power **proportional** to # of compromised TEEs

**“Design for Failure”**

**Open question:** How can TEE-using applications detect/mitigate effects of TEE-compromise?



Intel, [Hyperledger Sawtooth Documentation](#) (2015).

Chen et al., [“On Security Analysis of Proof-of-Elapsed-Time \(PoET\)”](#), SSS 2017.

# Summary

Concerns with applicability of hardware-supported TEEs remain

But compelling common-sense applications exist  
be **practical**; protect **end-users**; address **everyday** needs

Solutions must incorporate mitigations for:  
TEE **unavailability** or **compromise**

Design for failure

application- or system-level mitigations possible



<https://ssg.aalto.fi/research/projects/bcon/>  
BCon project, Academy of Finland



<http://www.icri-cars.org/>  
ICRI-CARS, Intel

# On dealing with TEE compromise

**Two types of settings where TEEs are useful:**

1. Improving **functionality** without compromising security: e.g., PoET
2. Improving **security** (esp. where none exists today): e.g., SafeKeeper

**TEE compromise is a major concern in Type 1 settings**

**In Type 2 settings, TEE compromise implies returning to current situation**