MOZI: Threshold FHE at Scale

Efficient and Resilient Threshold FHE for Large-Scale MPC

Presented by Yijia Chang @ The Hong Kong Polytechnic University

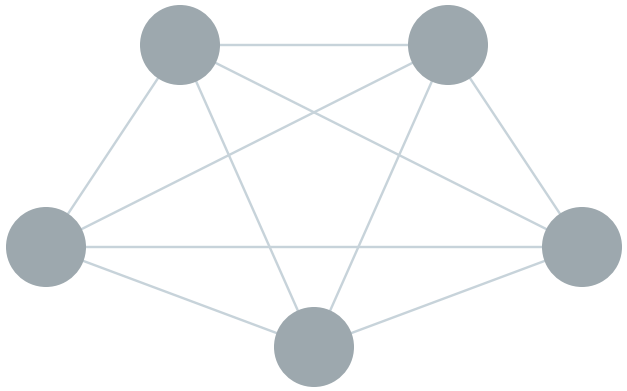
Team: Man Ho Au · Yijia Chang · Rongmao Chen · Xinyi Huang · Xianhui Lu · Moti Yung

Presented on 2026-Sep-30 in the virtual event TCPT3 (Threshold Call Preview Talks #3)

Why Large-Scale MPC?

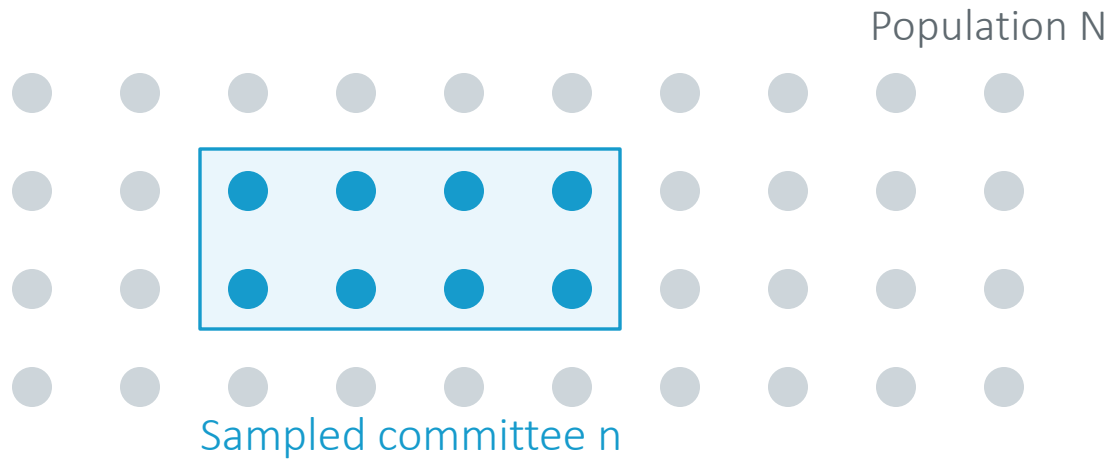
Large-scale MPC can distribute trust beyond a small, fixed committee

Small, fixed committee



The same few operators

Broader participant population



YOSO (You Only Speak Once)

One message per role, then retire

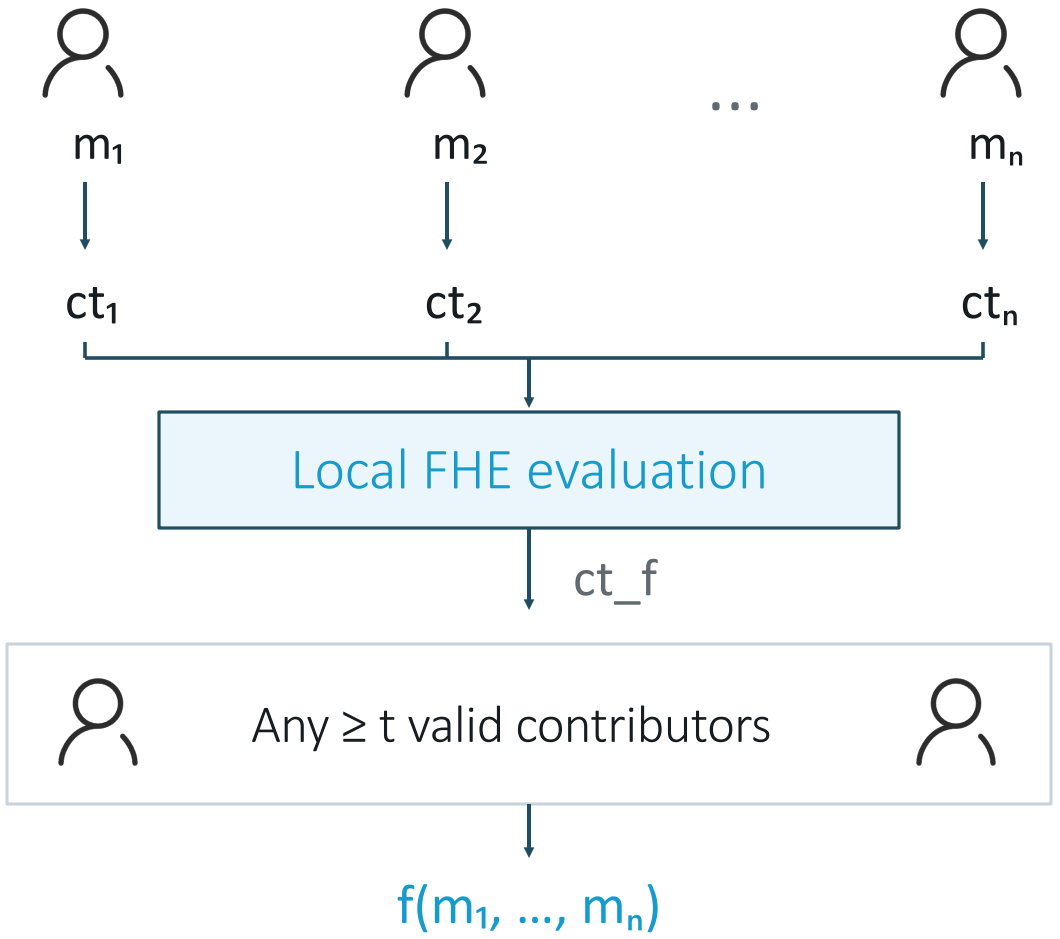
Threshold FHE for Large-Scale MPC

Threshold FHE avoids frequent party-to-party interactions for every computation step

1 Encryption

2 Evaluation

3 Threshold decryption



No joint computation

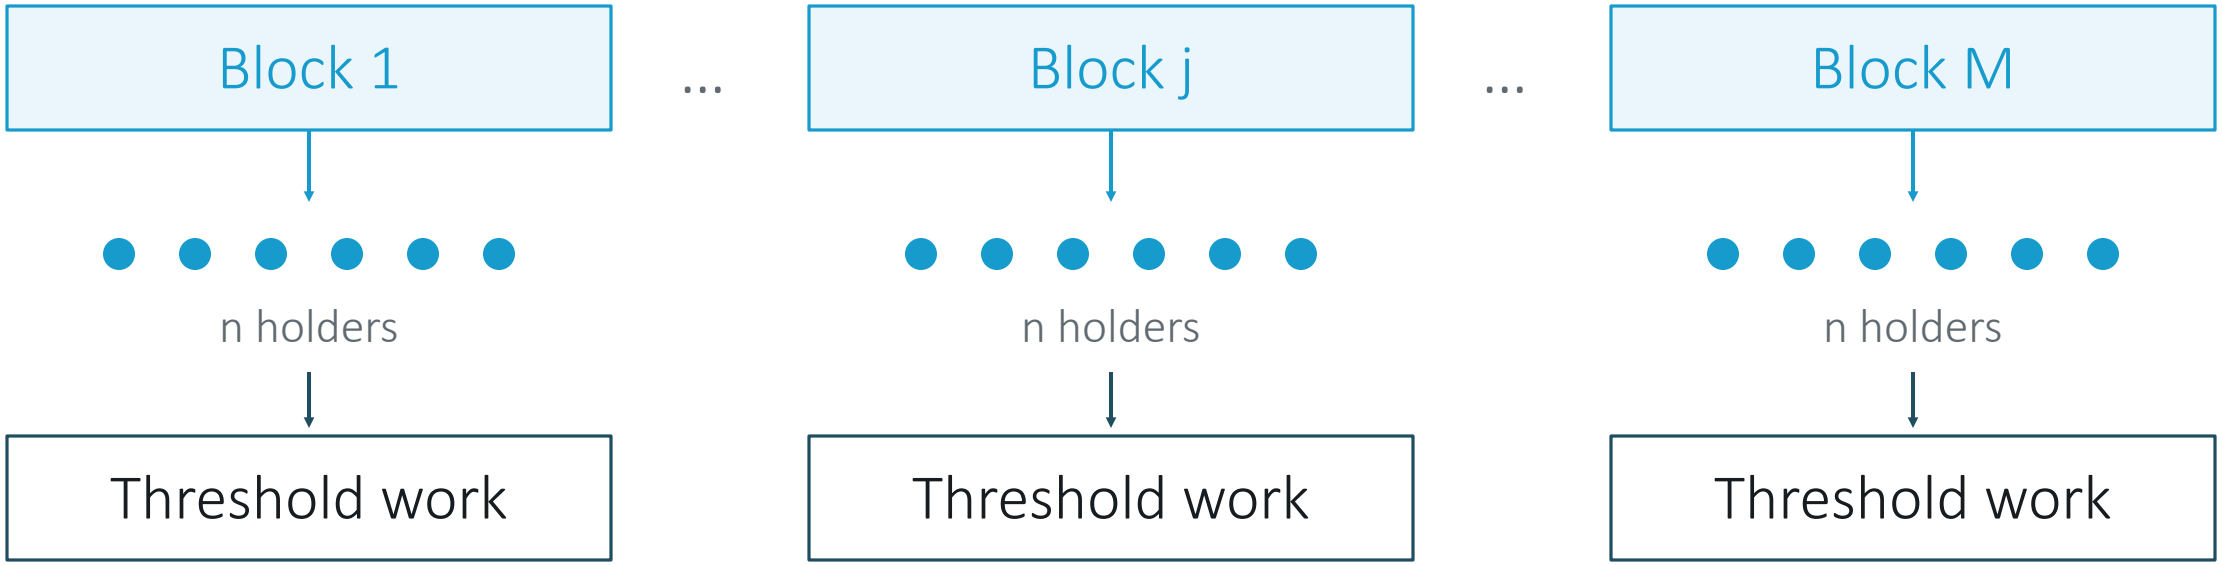
One evaluated ciphertext

Potential Bottleneck

Challenge 1: Low Complexity with # of parties and message size

Threshold decryption scales poorly when the number of parties and message size are large

A message of size M

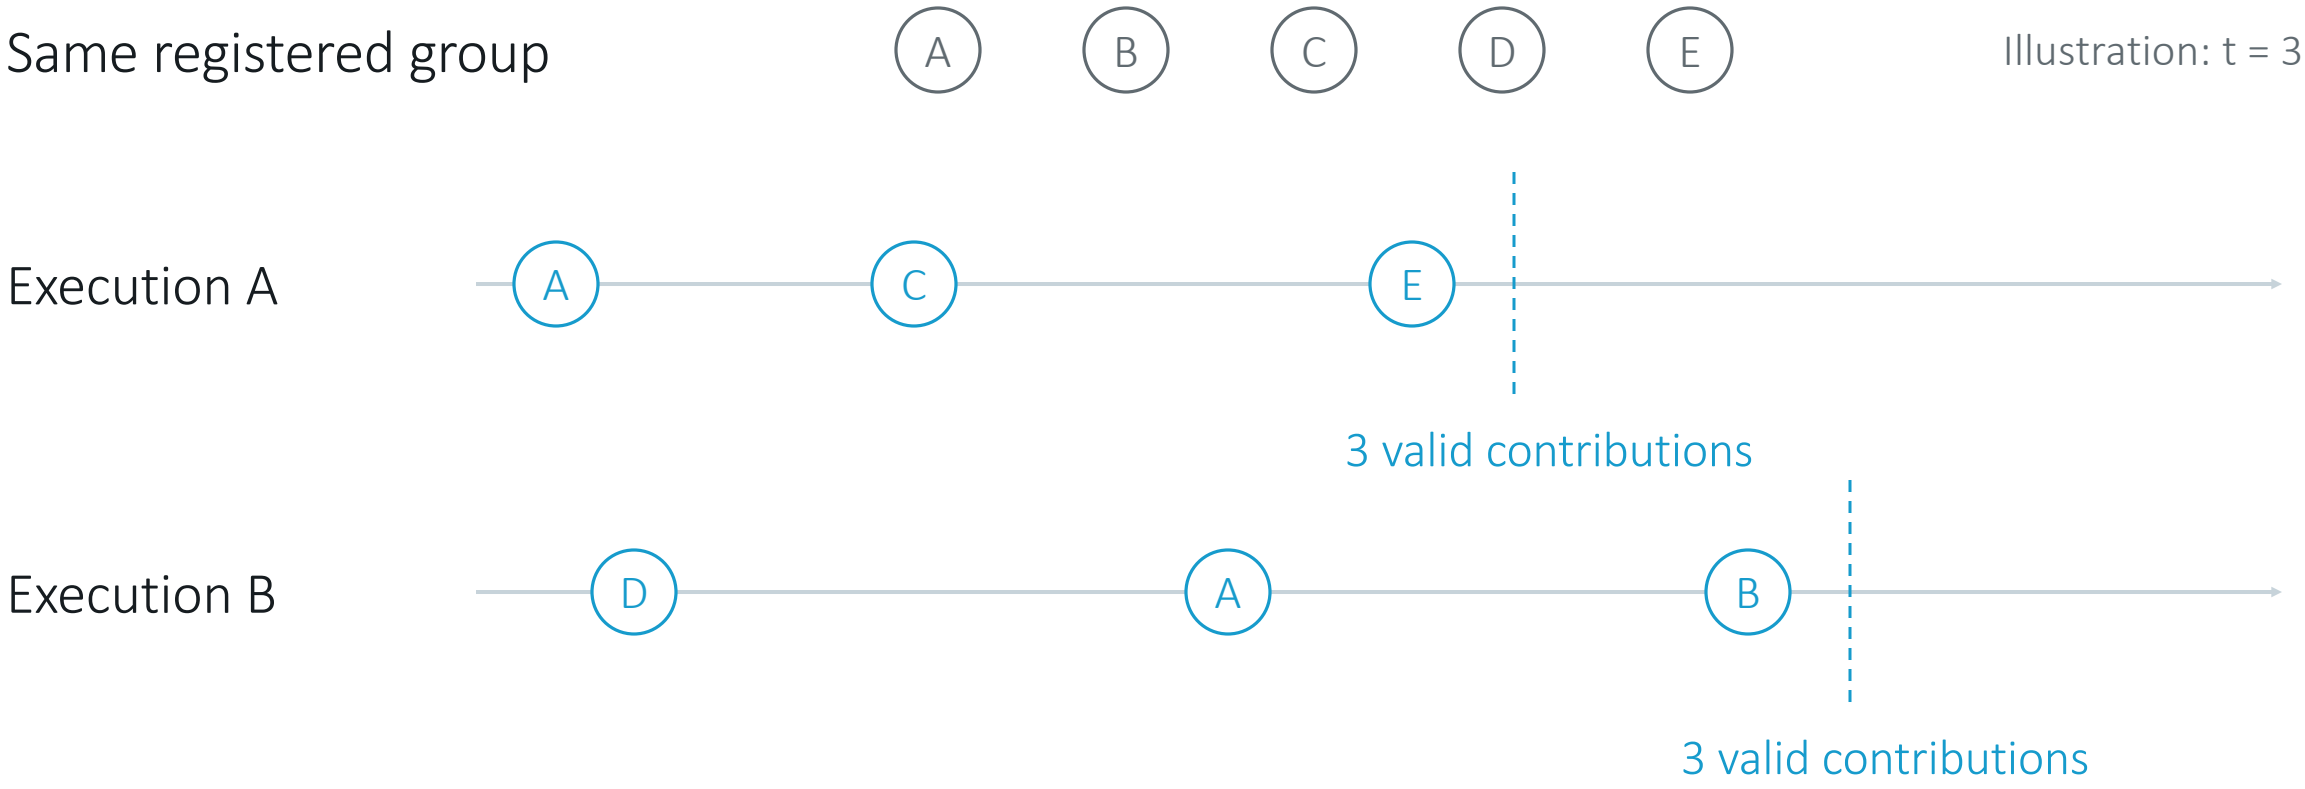


$$O(n^2) \times M = O(n^2 M)$$

Expensive operations × Long message

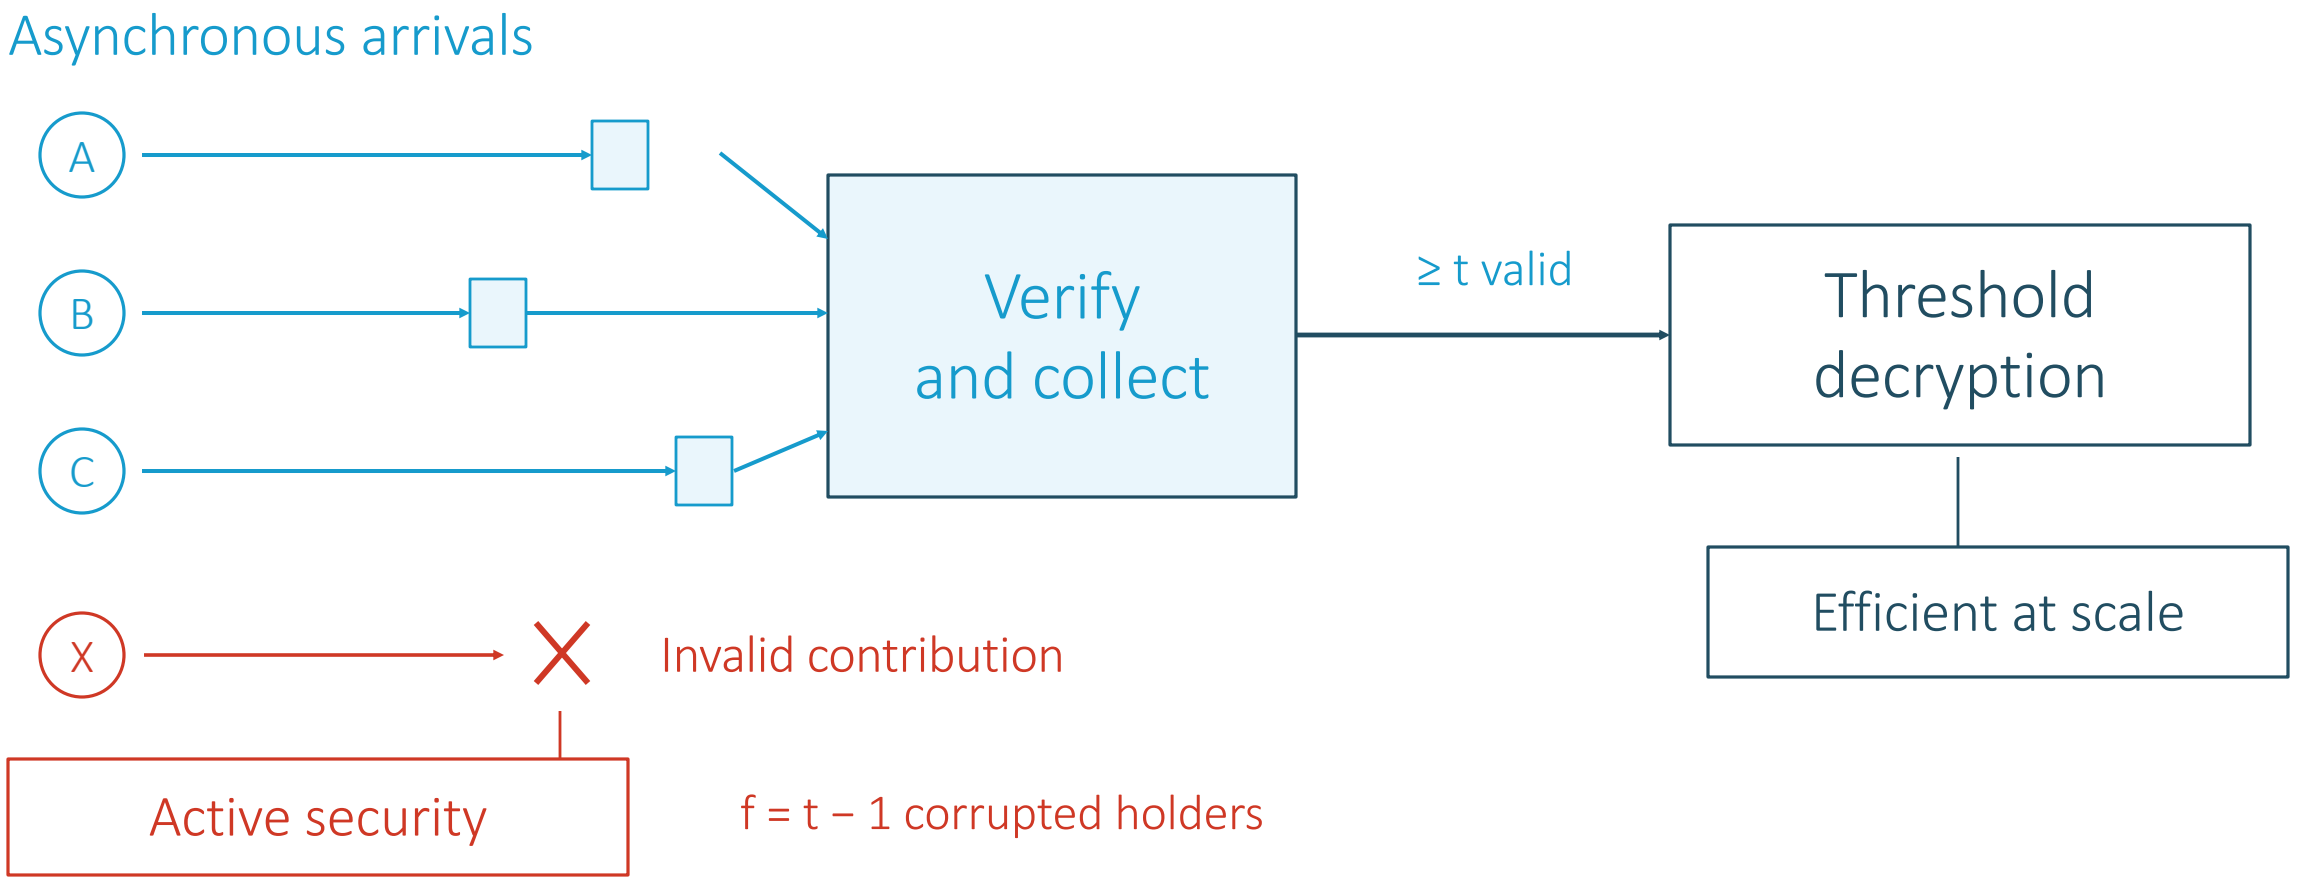
Challenge 2: An Unpredictable Contributor Set

In an asynchronous network with large scale, we cannot know the decryptor set in advance



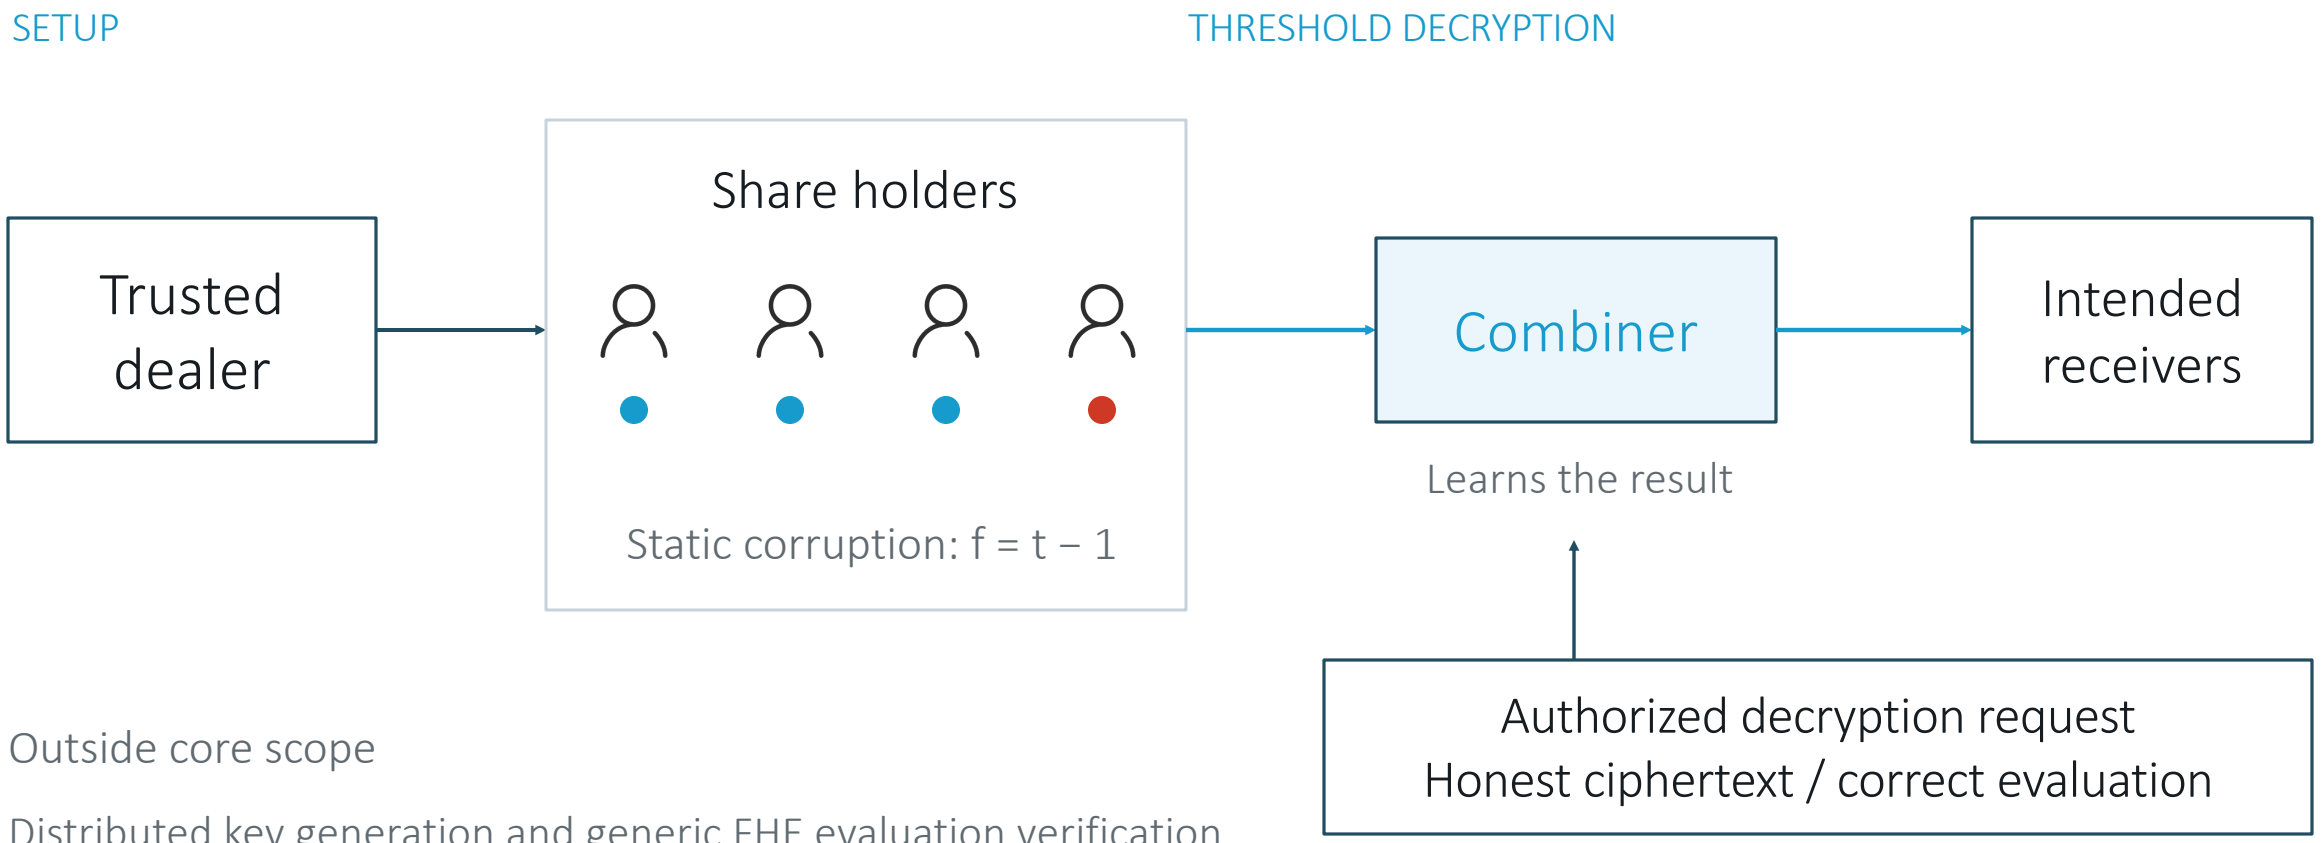
MOZI Design Goals

MOZI targets efficient, asynchronous threshold decryption with active security



System Model and Scope

MOZI assumes trusted setup, static adversary, and authorized decryption



Outside core scope

Distributed key generation and generic FHE evaluation verification

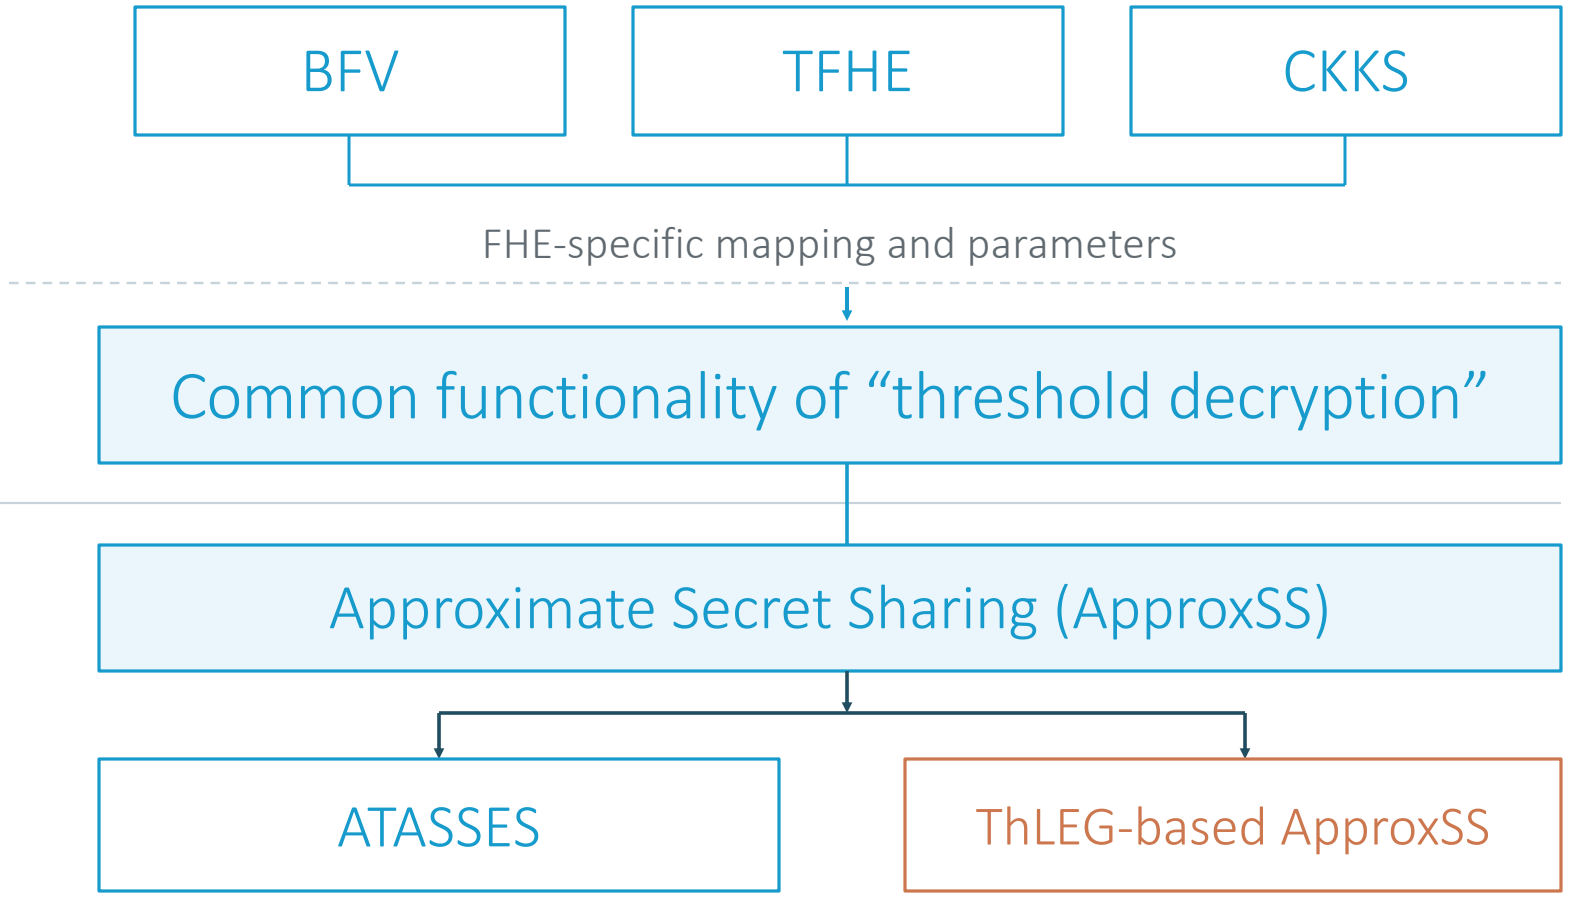
Two Design Levels for FHE Agility

Separating FHE from threshold decryption enables agility across FHE schemes.

Level 1

FHE agility

Change the FHE scheme;
reuse threshold decryption.



Level 2

Realization choice

Change assumptions
and reconstruction cost.

ApproxSS: The Core Design Problem

ApproxSS moves threshold-decryption design to approximate reconstruction.

Generic threshold-decryption construction

Local shares of b = $\Delta \cdot \mu + e$

Core design focus



≥ t holders

(< t shares hide b)

Privacy Hide the exact b

Only the permitted noisy value is revealed.

Correctness Bound the added noise η

Stay within the FHE decoding margin.

Noisy Shares: The Cost–Resilience Tension

Noisy shares couple the noise budget to the eventual reconstruction set.



Noise added to each share

The same weights also scale the noise.

$\lambda_i(S)$ can be large and depends on who participates.

Support arbitrary sets

Larger parameters / coordinated noise

Higher reconstruction cost

Fix the set in advance

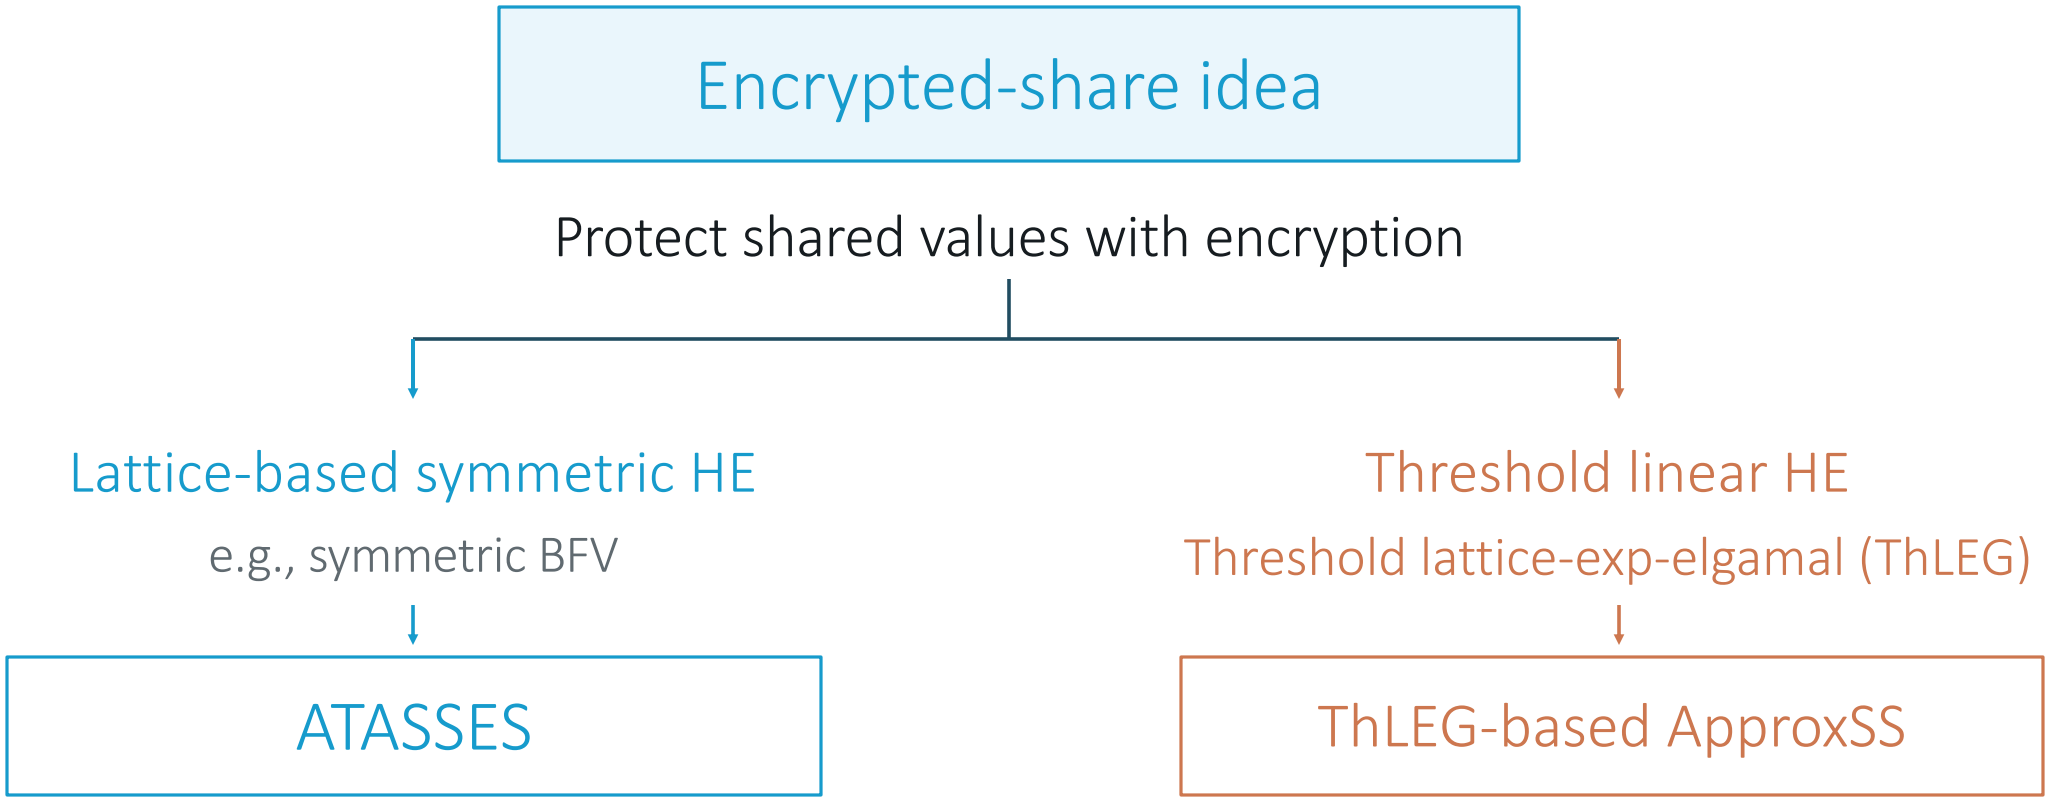
Depend on those particular holders

Less flexible asynchronous participation

Representative Shamir-based noisy-share designs (Chang and Li, USENIX Security 2025, §6.1).

Encrypted Shares: One Idea, Two Constructions

Encrypted shares combine lower cost with flexible participation.



Lower reconstruction cost + flexible qualified sets

Threshold Lattice-Exp-ElGamal (ThLEG)

ThLEG combines low-cost decryption and flexible participation for other threshold primitives

ThLEG

Threshold linearly homomorphic encryption

Low decryption cost

O(n log n + nD + M)

D: adapter key size; M: message size

High participation resilience

Any qualified set of ≥ t holders

No preselected set; asynchronous arrivals

Use in MOZI

ApproxSS → threshold FHE

Potential reuse

Threshold signatures

CL-based ECDSA / BBS+: Wong et al., NDSS 2024. ThLEG adaptation requires protocol-specific proofs.

Fang and Yuan: Cost and Assumptions

Fang targets post-quantum security, while Yuan reduces the threshold overhead.

方

MOZI-Fang

Square / lattice foundation

圆

MOZI-Yuan

Circle / curve component

Compared noisy-share baselines: $O(n^2M)$

$O(n^2 + nM)$

$O(n \log n + nM)$

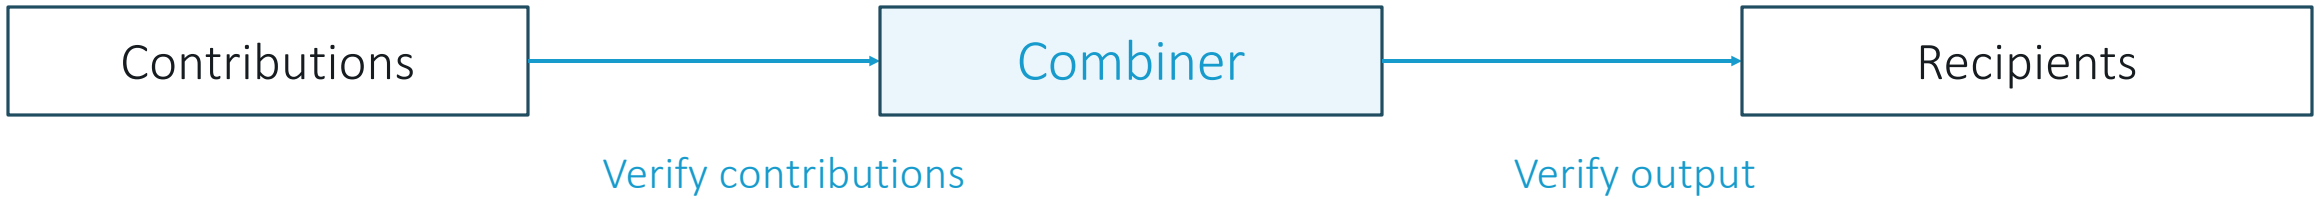
Realization	ATASSES	ThLEG-based ApproxSS
Target	Post-quantum security	Lower threshold overhead
Tradeoff	Encryption and key-material cost	Quantum-vulnerable assumptions

n: holders; M: message size; t = Θ(n). Fixed parameters, no preprocessing.

Combiner + one holder. Yuan bound derived with noise aggregation. MOZI proof/network overhead excluded.

Security and Output Delivery

The combiner's behavior determines which output-delivery guarantee we target.



Semi-honest combiner	$t + f \le n$	Guaranteed output delivery
Malicious combiner	$1 \le t \le n$	Security with abort

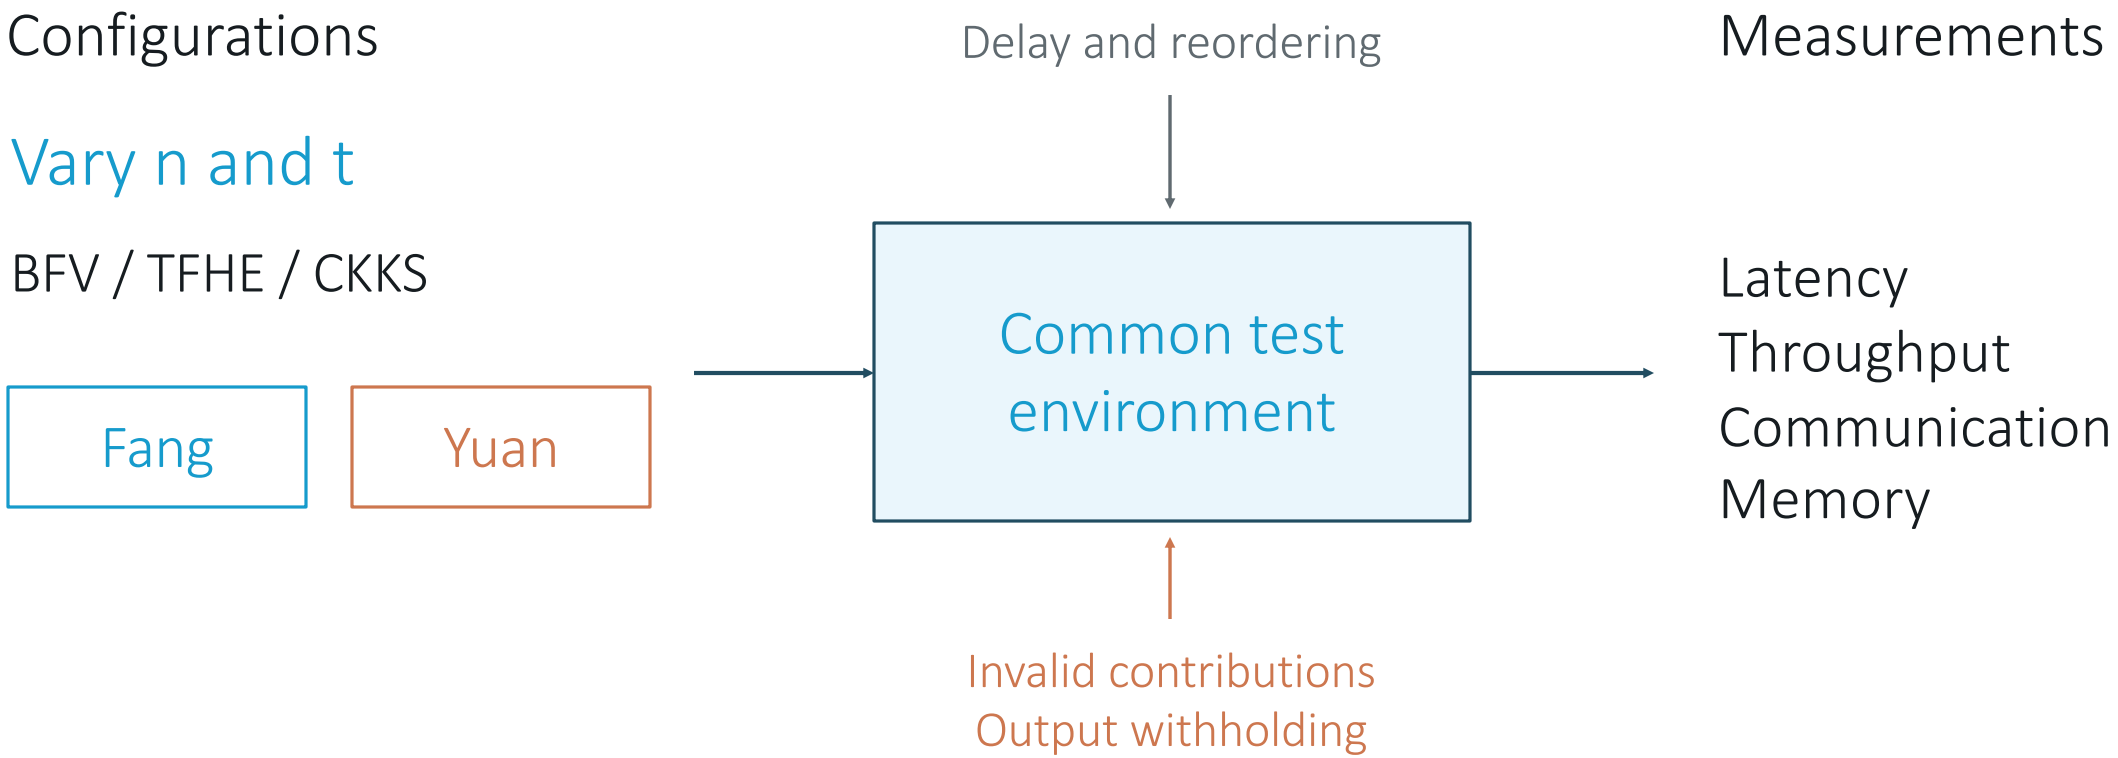
Both profiles use $f = t - 1$

Parameter targets	FHE: 128-bit classical	ThLEG: ~128-bit classical

Ongoing work — Malicious-combiner GOD for $f < n/3$

Evaluation Plan

We will test threshold decryption cost and resilience, using Beaver triples as an example.

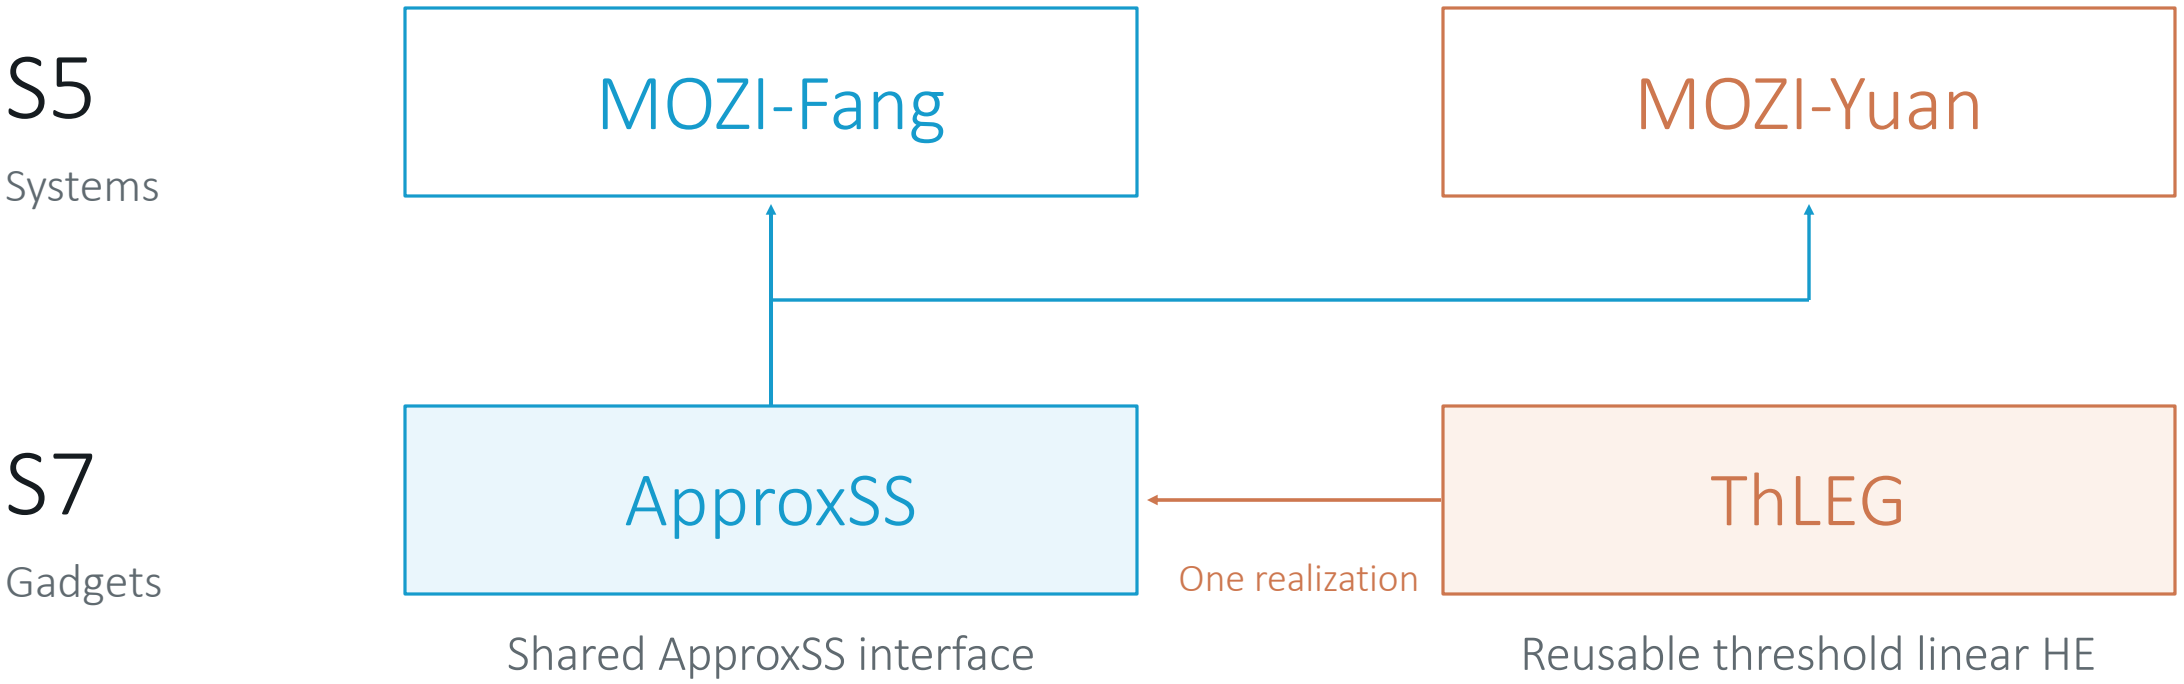


MPC workload example

Usable authenticated Beaver triples after checking and sacrifice

Planned MOZI Submission

MOZI plans two threshold FHE systems and two reusable cryptographic gadgets.



Specifications, reference code, tests and benchmark scripts

Thank you!

Any questions?