MPI-SGX: Enabling Confidential Computing for MPI Parallel Applications with Intel SGX Technology

*¶**Kota Shimojima** *Hayato Yamaki *Hiroki Honda †Shi'nichiro Matsuo ‡¶Atsuko Takefusa *¶Shinobu Miwa

* The University of Electro-Communications, † Georgetown University, ‡ National Institute of Informatics, ¶RIKEN R-CCS

Introduction

- The demand for confidential HPC is growing.
- Trusted Execution Environments (TEEs) have gained attention in HPC due to their ability to securely execute code with minimal performance overhead.
  ※TEEs such as Intel Secure Guard Extension (SGX) offer a hardware-assisted sandbox.
- No existing study presents high performance parallel confidential computing over multiple TEE-enabled compute nodes.

Contribution

- We propose the first SGX-based parallel computing system.
- We develop a new secure MPI library called MPI-SGX.

Threat Model

Assume two adversaries in parallel computing systems

1. **An intra-node adversary** can control the OS, hypervisor, and processes on any node and can read data in use and in memory.
2. **An inter-node adversary** can intercept, modify, replay, and drop messages on the interconnect (as in the Dolev–Yao model).

We do not consider threats unrelated to information leakage or system privileges.

SGX-Based Parallel Computing System

Propose a novel parallel computing system based on SGX.

1. Each SGX offers a secure computing environment for an application process to protect computation data against intra-node adversaries.
2. The key management system running on the login node provides an authorized user with an encryption key needed to access sensitive data.
3. Communication between the compute nodes and the login node is encrypted, thereby protecting communication data against inter-node adversaries.
4. The verification server is used to check the integrity of the SGXs used for computation.

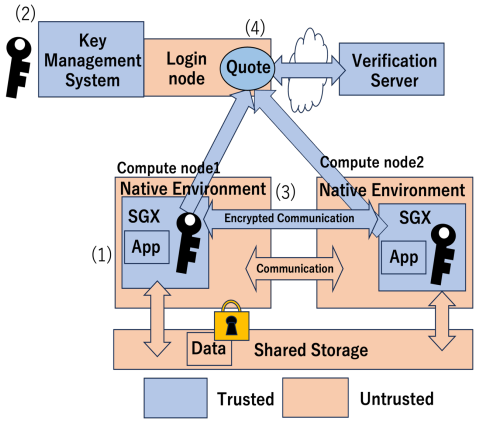

[Proposed system]

MPI-SGX

- Our proposed MPI-SGX is an MPI communication library that provides enabling secure inter-node channels for parallel applications running on SGX + Gramine.
  ※ Gramine is a lightweight library OS for executing unmodified Linux binaries on SGX.
- MPI-SGX adds the following three functions into the original MPI library, so users can execute their applications confidentially on multiple SGXs by simply linking our library.
  - **Encrypted-Communication:** Use AES-GCM via OpenSSL for efficient symmetric encryption, reducing performance overhead compared to RSA.
  - **Key Distribution**: The root process generates a random symmetric key, encrypts it with each receiver's public key, and then sends it to other processes.
  - **Remote Attestation**: Based on DCAP, performed at initialization and periodically during execution. All nodes must pass verification. Otherwise, the execution is aborted.

We implement MPI-SGX on Open MPI 1.5.0 for compatibility with Gramine+SGX.

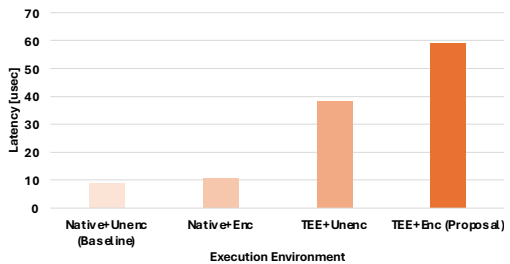

[PingPong- Latency]

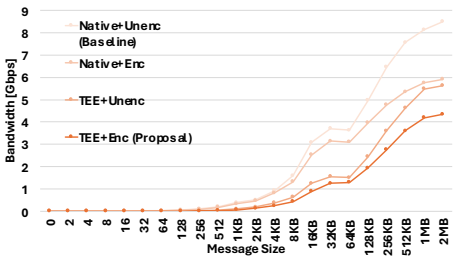

[PingPong-Bandwidth]

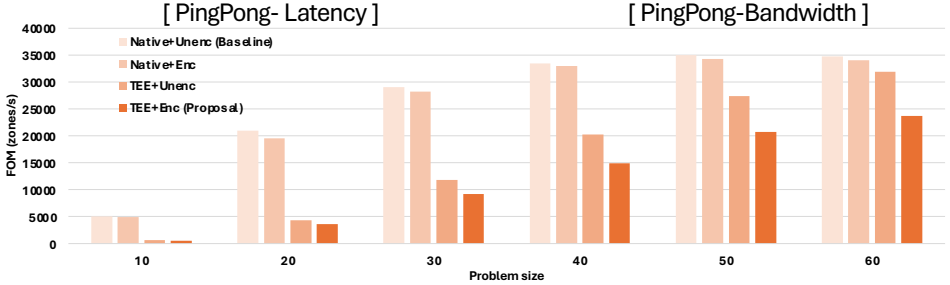

[LULESH Results]

Experiments

- We select PingPong and LULESH and compare the perfomances.
  - TEE+Enc shows a 6.6x larger latency than Native+Unenc.
  - TEE+Enc degrades the bandwidth by up to 49% (on a message of 2MB) compared to Native+Unenc.
  - In LULESH, the performance degradation of TEE+Enc decreases as the problem size increases.

[Experimental settings]

Name	Remarks
CPU	1x Intel(R) Xeon(R) Gold 5317 CPU, 3.0GHz,12C24T
Memory	64GiB DDR4-3200 (32GiB enclave memory)
Network	10G Ethernet
OS	Ubuntu 20.04.6 LTS

- The performance overhead of MPI-SGX is several orders of magnitude smaller than that of the existing confidential computing techniques, e.g., homomorphic encryption.
- We will make further efforts to improve the performance of MPI-SGX.

Summary

- We proposed the SGX-based parallel computing system and MPI-SGX to realize confidential computing on supercomputers.
- Our future work will perform a comprehensive evaluation with real HPC applications.

This work was supported by JST, PRESTO Grant Number JPMJPR22P9.0