Gu JY, Li H, Xia YB *et al.* Unified enclave abstraction and secure enclave migration on heterogeneous security architectures. JOURNAL OF COMPUTER SCIENCE AND TECHNOLOGY 37(2): 468–486 Mar. 2022. DOI 10.1007/s11390-021-1083-8

Unified Enclave Abstraction and Secure Enclave Migration on Heterogeneous Security Architectures

Jin-Yu Gu[1,2] (古金宇), Hao Li[1,2] (李　浩)
Yu-Bin Xia[1,2,*] (夏虞斌), *Senior Member*, *CCF*, *Member*, *ACM*, *IEEE*
Hai-Bo Chen[1,2] (陈海波), *Distinguished Member*, *CCF*, *ACM*, Cheng-Gang Qin[3] (秦承刚), and
Zheng-Yu He[3] (何征宇)

[1] *Engineering Research Center for Domain-Specific Operating Systems, Ministry of Education, Shanghai 200240, China*
[2] *Institute of Parallel and Distributed Systems, Shanghai Jiao Tong University, Shanghai 200240, China*
[3] *Ant Group, Hangzhou 310099, China*

E-mail: {gujinyu, lihao, xiayubin, haibochen}@sjtu.edu.cn; {chenggang.qcg, zhengyu.he}@antgroup.com

Received October 18, 2020; accepted February 21, 2021.

Abstract Nowadays, application migration becomes more and more attractive. For example, it can make computation closer to data sources or make service closer to end-users, which may significantly decrease latency in edge computing. Yet, migrating applications among servers that are controlled by different platform owners raises security issues. We leverage hardware-secured trusted execution environment (TEE, aka., enclave) technologies, such as Intel SGX, AMD SEV, and ARM TrustZone, for protecting critical computations on untrusted servers. However, these hardware TEEs propose non-uniform programming abstractions and are based on heterogeneous architectures, which not only forces programmers to develop secure applications targeting some specific abstraction but also hinders the migration of protected applications. Therefore, we propose UniTEE which gives a unified enclave programming abstraction across the above three hardware TEEs by using a microkernel-based design and enables the secure enclave migration by integrating heterogeneous migration techniques. We have implemented the prototype on real machines. The evaluation results show the migration support incurs nearly-zero runtime overhead and the migration procedure is also efficient.

Keywords heterogeneous trusted execution environment (TEE), enclave abstraction, enclave migration

1 Introduction

As an emerging computing paradigm, edge computing [1–4] has gained more attention in recent years because it allows services to become closer to clients or data production sources. Owing to the promising feature of "close-to-data/client", edge computing can significantly reduce network communication cost and thus bring better quality of service, e.g., extremely low latency for requests. Nowadays, it has been used in plentiful application domains, such as computation offloading from cloud to smart home or city [5, 6], real-time analytics [7, 8], and so on, to address the concerns of response time requirement and bandwidth cost limitation.

The mobility of clients (e.g., mobile users) and data sources (e.g., intelligent vehicles) makes runtime service migration become an indispensable requirement in edge computing [9–12]. With migration, a service can keep running on the edge server nearest to the client, which may change from time to time, in order to keep latency low. Otherwise, dramatic performance degradation may occur, and qualified service continuity is difficult to ensure. In addition, migration is also important for meeting other demands of edge computing, like relieving congested edge servers and leaving servers

Regular Paper

This work is supported in part by the National Key Research and Development Program of China under Grant No. 2020AAA0108502, the National Natural Science Foundation of China under Grant Nos. 61972244, U19A2060, and 61925206, and the HighTech Support Program from Shanghai Committee of Science and Technology under Grant No. 19511121100.

*Corresponding Author

that may fail (e.g., before running out of battery).

However, research on security issues of service migration in edge paradigms is still nascent and limited [9, 13, 14]. Compared with traditional cloud computing, there are two main security-related differences when edge servers involve. First, the owners of the edge servers may be different from the cloud providers and could be curious or even malicious. Thus, the providers of the service applications (e.g., the application developers) have concern about their intellectual property which may be easily stolen by the edge server owner who controls the physical machine as well as the whole system software stack. The end users also cannot ensure the service application is correctly running on the edge servers. Second, the edge servers are easier to be attacked compared with cloud servers because cloud servers usually face remote attacks only while attackers are easier to physically access edge servers and thus have more attack means (e.g., conducting physical attacks). Such differences are obstacles for migrating services among edge servers as well as from cloud to edge.

In this paper, we propose to utilize the hardware-assisted trusted execution environment (TEE) to mitigate the above security threats and enable secure service migration. TEE is suitable for protecting private code and data on untrusted platforms. For example, Intel SGX has been adopted by some major cloud providers and ARM TrustZone has been well-used on smartphones. Generally speaking, when accommodated in a hardware TEE, a benign service application can protect itself and users' input from malicious software, including OS and compromised peripherals. Nevertheless, edge servers can deploy CPUs from different vendors, such as Intel, AMD and Huawei, which, inherently, means that their equipped TEEs are heterogeneous, like Intel SGX [15], AMD SEV [16], and ARM TrustZone [17].

Therefore, we propose UniTEE which gives a unified TEE abstraction for hiding the hardware heterogeneity from applications. UniTEE adopts the programming model of SGX applications for its flexibility and popularity. Specifically, an application can partition itself into secure-(in)sensitive parts and build one or more hardware-secured TEEs (named enclaves) to run the secure-sensitive ones. An enclave can offer strong guarantees of both confidentiality and integrity for the secure code/data inside despite being executed in an untrusted environment, which can be extremely suitable for outsourced computation [18–21]. No matter what the underlying hardware TEEs are, UniTEE provides unified programming APIs including creating, attesting, invoking, and destroying enclaves. As AMD SEV and ARM TrustZone do not provide enclaves like Intel SGX, we leverage hardware-software co-designs for building SGX-like enclaves on those platforms. AMD SEV uses virtual machine (VM) as the granularity of its TEE and supports concurrently running at most 15 secure VMs, which does not fit the programming model of UniTEE. Therefore, we deploy a trusted microkernel in the supervisor mode of a secure VM and then let the microkernel to build user-level isolated enclaves. An application can construct its enclaves in the secure VM by sending requests to the trusted microkernel. ARM TrustZone enables the CPU to have two modes named normal world and secure world, respectively. UniTEE achieves the same enclave abstraction by deploying the trusted microkernel in the secure world to be the enclave manager. Thus, by combining the tiny software layer (the trusted microkernel) and the hardware TEE (either a secure VM of SEV or the secure world of TrustZone), UniTEE provides SGX-like enclave abstractions and thus unifies the TEE programming model on Intel SGX, AMD SEV, and ARM Trustzone. Besides, for easing programming, it provides an enclave-management library for an application to control its enclaves' life cycle, including creation, attestation, interaction and deletion. It also provides a C library (based on musl-libc) to ease the development and deployment of in-enclave code, as well as to be compatible with legacy code.

A unified enclave abstraction enables programmers to develop secure applications without considering the differences of the underlying TEEs. Nevertheless, it is not enough for migrating applications between edge nodes because heterogeneous TEEs use different instruction set architectures (ISAs). Therefore, UniTEE further integrates heterogeneous-ISA migration techniques [11, 22, 23] to hide the heterogeneity of enclave ISAs and support enclave migration [24] at runtime. The enclave code will be compiled into different binaries for different ISAs, but every symbol (a variable and a function) has the same offset in different binaries. No matter on which architecture, these symbols will always be loaded at the same virtual addresses at runtime, which significantly simplifies the (cross-architecture) migration procedure because the pointers to them will still be valid after migration. For migrating an enclave running on the source machine, the target machine will first launch a virgin enclave with the binary for its architecture and then receive and restore the enclave checkpoint (memory data and execution context) from the

source machine. For ensuring security, the checkpoint generation should not rely on the untrusted software including the OS. Thus, an enclave on the source machine will generate a consistent checkpoint by itself. Specifically, UniTEE adds a control thread in each enclave as a part of the framework. After receiving a migration request, this thread will wait for all the enclave threads to enter a quiescent state and then make a checkpoint by encrypting and dumping the enclave states. The encryption key is negotiated by the source enclave and the target enclave and it will protect both the confidentiality and the integrity of the checkpoint during the transfer process.

UniTEE provides a software development kit (SDK) for programmers, and they can develop secure edge applications without awareness of the underlying TEE hardware or the migration mechanisms. We present a prototype implementation and evaluation on an Intel (Skylake i7-7700) machine, an AMD (EPYC 7281) machine, and an ARM (HiKey970) machine, respectively. The evaluation results show: 1) our SDK can support many real-world applications and the migration mechanism incurs negligible overhead; and 2) the latency of heterogeneous enclave migration is acceptable and mainly decided by the network latency.

In summary, this paper makes the following contributions:

• a unified enclave abstraction on Intel SGX, AMD SEV, and ARM TrustZone exposed by UniTEE;

• a design of secure enclave migration between heterogeneous TEEs enabled in UniTEE;

• a real implementation and evaluation of UniTEE.

2 Motivation and Background

2.1 Motivation

Application (Service) Migration Is One of the Most Critical Features in Edge Computing[8–11, 25–27]. First, migration can relieve congested edge servers and thus achieve better load balance. Second, migration can bring better fault tolerance (e.g., migrating applications that run on a low-battery edge machine) and ease the edge server upgrading (e.g., migrating running applications to other servers first and then updating an idle machine). Third, migration is important to ensure Quality-of-Service (QoS) in edge computing because of the high mobility of client devices such as smartphones or intelligent cars. Specifically, the latency between a client and an edge service may vary because of the client's mobility, which can impact the overall performance, i.e., the service quality. Fig. 1 presents such an example which shows the latency impact on Vedis, a popular key-value store on edge. If latency-sensitive services can be migrated to follow clients, they can show much better performance.

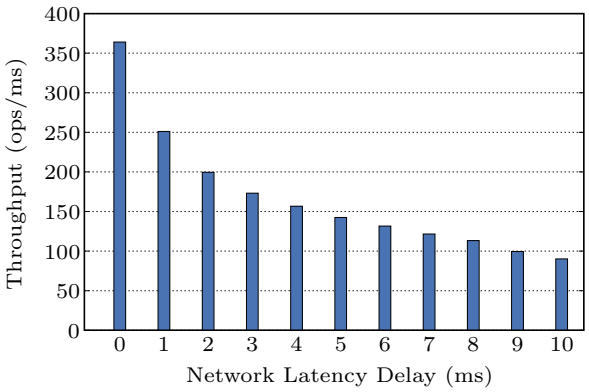

Fig. 1. Throughput of the Vedis service that runs on an Intel Skylake machine. The client device is an ARM HiKey board. The network latency varies from 1 ms to 10 ms (typical latencies in edge computing[11, 28]).

Security Concerns on Migration for Both Service Providers and Clients. Although service providers have the above motivations to migrate their services between different servers, security concerns may force them to abandon migration. Different from cloud servers, which are aggregated together in data centers and managed by the same cloud provider, edge servers are more disaggregated and can be managed by different owners. If one service application is allowed to migrate from cloud to edge and between different edge nodes, the service provider faces the risk of leaking digital property because the owner of some edge server may be curious or even malicious. The owner has the full control of the edge server and can deploy malicious system software (e.g., OS and hypervisor) or compromise them. If a service runs on such a server, the server owner can easily retrieve all the code, data, and runtime states of the service, which means the loss of the digital property to the service provider. Moreover, when the service is controlled by an untrusted server owner, the clients of the service also worry about the security: the client data sent to the service may be stolen, or the service may not faithfully handle the requests.

Hardware TEEs Bring a Potential to Solve the Security Problems. Nowadays, hardware support for secure computing, i.e., Intel SGX, ARM TrustZone, and AMD SEV, gains more and more attention in both the academic and the industry area. These hardware security extensions can protect security-sensitive applications from attackers through providing a hardware-

secured trusted execution environment (TEE). A TEE can shield an application's code and data from external accesses by other software, including higher-privileged software like OS and hypervisor. Besides memory protection, it can also provide tamper-resistant execution for the protected application. Therefore, the hardware-supported TEE technology is a promising candidate for protecting applications in untrusted cloud/edge servers where the entire software stack and the infrastructure owner are not trustworthy.

However, the Heterogeneity of Servers in Edge Computing Leads to Two-Fold Challenges. In terms of application programming, different servers are equipped with different kinds of TEEs which give heterogeneous programming abstractions (Challenge-1). Writing code for every abstraction not only makes the application development inefficient but also brings difficulties to runtime migration. In addition, heterogeneous hardware TEEs make the migration procedure of protected applications challenging (Challenge-2) for two reasons. First, they use architecture-specific instructions, registers, etc., which are different from each other. Second, they cannot be accessed by privileged system software such as OS and hypervisor which play important roles in traditional migration (e.g., OS will stop an application and then send its memory data).

In this paper, we make an attempt to solve the above two challenges on how to program security-sensitive applications with different heterogeneous TEEs and how to migrate them between the heterogeneous TEEs. Besides in edge computing as mentioned above, our work may also be used in cloud computing where heterogeneous hardware TEEs and migration are also required [22, 24, 29–33], for example, secure cross-cloud migration is needed in joint cloud computing [34].

2.2 Background of Hardware-Secured TEEs

Intel SGX. Intel SGX [15] can protect user-level computations through providing a hardware-secured execution environment called an enclave. An enclave's memory pages reside in the EPC (Enclave Page Cache) which is a part of the memory region that will be automatically protected by the CPU. Although the hypervisor and OS retain their ability to manage EPC memory (e.g., swap EPC page), they cannot break the memory data's confidentiality and integrity. Moreover, Intel SGX also provides tamper-resistant execution to enclaves and enables remote attestation, which means that an enclave can prove its identity to a remote party. Thus, researchers have proposed to leverage SGX to protect outsourced applications [19, 21, 35–38] and cloud vendors have started to explore the commercial usage of SGX [39]. Specifically, an application can be separated into trusted and untrusted parts, and the trust parts can be executed in one or more enclaves. An SGX enclave resides in the address space of its host application while its memory can only be accessed by itself. A thread has to enter an enclave through executing an EENTER instruction and exit from the enclave with an EEXIT instruction. Moreover, the CPU can help an enclave to produce a verifiable proof that identifies its memory contents. A remote party, e.g., the enclave owner, can leverage official attestation services like Intel Attestation Service (IAS) to assess the trustworthiness of the proof. Such a procedure is called SGX remote attestation.

AMD SEV. AMD proposed SEV [16] to protect outsourced computing on untrusted servers, whose support has been integrated into existing system software stacks. Different from Intel SGX, which can build TEEs inside applications, the granularity of a TEE in SEV is a secure virtual machine (VM). Tenants can run their applications inside a secure VM which is protected as a whole by the SEV hardware. SEV supports at most 15 secure VMs, and each of them has a unique identifier (ASID). Inside the CPU, all the secure memory of the VM is tagged with the VM's ASID, which prevents the memory content from being accessed by anyone other than the owner VM. When the secure memory data leaves/enters the CPU, it is automatically encrypted/decrypted by the memory controller with a key bound to its owner VM. These keys are managed by a secure co-processor and will never be exposed. A secure VM can decide whether a memory page is secure by setting one bit (C-bit) in the corresponding guest page table entry. Once the bit is set, the CPU will treat the memory page as secure and then protect it transparently. Otherwise, the access to the memory page is not restricted, i.e., the page can be accessed by the hypervisor. Besides memory protection, SEV also protects a secure VM's execution states during runtime. Recently, AMD has proposed further extensions named SEV-SNP① which helps to mitigate the memory integrity problems of SEV [40, 41]. Therefore, even in the face of compromised privileged software, SEV can also

①SEV Secure Nested Paging. https://www.amd.com/system/files/TechDocs/56860.pdf, Feb. 2021.

protect both the confidentiality and the integrity of its TEEs.

ARM TrustZone. ARM proposes the TrustZone [42] technology as its hardware security extension since ARMv6 architecture. With TrustZone, the CPU has two execution environments, named normal world and secure world, separately. Both worlds have their own user space and kernel space, while the latter is used as the TEE on ARM. Usually, a commodity OS and non-security-sensitive applications run in the normal world while a secure OS (e.g., OPTEE②) and security-sensitive applications run in the secure world. The two worlds can switch to each other through the highest privilege mode, monitor mode. One world can execute an SMC (Secure Monitor Call) instruction to trap into the monitor mode, and then a secure monitor in the monitor mode helps to finish the world switch. TrustZone can also partition all the physical memory resources into the normal part and the secure part, and ensure that the normal world cannot access the secure memory part while the secure world can access the entire memory. Thus, two worlds can exchange data through the normal memory part. Moreover, the secure world can adjust the memory resource partition according to runtime requirements. As TrustZone is widely deployed in ARM platforms such as smartphones and tablets, it has already been used to protect security-critical computation and data [1, 43–47].

3 Unified TEE Programming Abstraction

In this section, we first give a brief analysis of the three commercial TEE abstractions, which will explain why UniTEE chooses the SGX-like abstraction as the unified one. Then, we describe how to achieve the unified TEE abstraction on different security hardware. Last, the main programming interfaces of such an abstraction will be introduced.

3.1 Abstraction Analysis

As shown in Fig.2, Intel SGX supports constructing multiple enclaves (fine-grained TEEs) inside an application, i.e., in the application's address space, which allows programmers to divide an application into one untrusted part and one or more trusted parts. The latter ones are used to protect security-sensitive code, data, and execution. There are several typical usages of SGX for application protection (different granularities). First, programmers can put an unmodified application together with a library OS into a single enclave (e.g., Graphene-SGX [21]) which may also be deployed as a guest VM on untrusted servers (e.g., Haven [35]). Second, programmers can run a container in an enclave to enhance security (e.g., SCONE [19]). Third, programmers can partition an application into mutual-distrusted parts manually [38] or automatically [48] and then utilize enclaves for isolation. In brief, SGX enclave abstraction not only allows a relatively unlimited number of TEEs, but also promises flexibility in the isolation granularity.

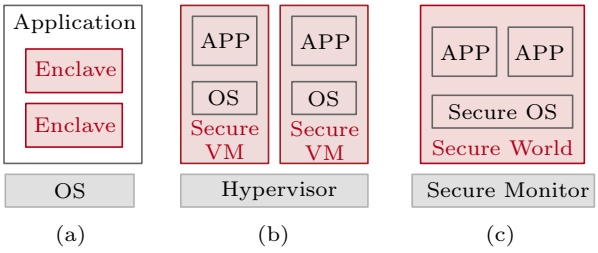

Fig.2. Three commercial TEE abstractions. (a) Intel SGX. (b) AMD SEV. (c) ARM TrustZone.

In contrast, AMD SEV supports at most 15 secure VMs as TEEs, which leads to two drawbacks: 1) the TEE number is too limited to accommodate different applications; 2) the TEE granularity is too coarse-grained to meet different requirements. Similarly, ARM TrustZone provides only-one secure world as TEE. Although prior studies proposed to multiplex the secure world by deploying a secure OS or using virtualization [31], they only considered protecting a whole application instead of fine-grained protection enabled by SGX.

Therefore, UniTEE embraces the flexible abstraction of SGX and allows to build SGX-like enclaves with any of the three hardware-security technologies. Programmers can develop secure applications against a unified abstraction without concerning the underlying hardware TEEs.

3.2 System Architecture

To provide SGX-like abstraction with AMD SEV or ARM TrustZone, the first problem to solve is that they cannot provide an unlimited number of hardware TEEs as enclaves. To this end, UniTEE deploys a security-oriented microkernel in one hardware TEE, i.e., the secure world of ARM TrustZone or a secure VM of

②https://github.com/OP-TEE/optee_os, Feb. 2021.

AMD SEV, and then leverages the microkernel to construct an unlimited number of software TEEs as enclaves. The microkernel runs in kernel mode while the enclaves managed by it are running in user mode.

As shown in Fig.3, the microkernel creates a new address space for building a new enclave, which is similar to a traditional process on the microkernel. Nevertheless, an enclave logically belongs to some application that runs in the normal VM on an SEV-capable machine or in the normal world on a TrustZone-capable machine. The trustworthy microkernel guarantees both the isolation between different enclaves and the isolation between an enclave and all the untrusted software in the normal VM or the normal world. Specifically, it assigns different enclaves with different page tables and thus achieves the memory isolation between them; it leverages the hardware-security mechanism to ensure the memory isolation between an enclave and the untrusted software, i.e., the enclave uses the secure memory (in the secure VM or the secure world) that cannot be accessed by the untrusted ones, including the privileged OS. Besides, the microkernel also manages the enclaves' runtime states (execution context) and ensures the states' confidentiality and integrity.

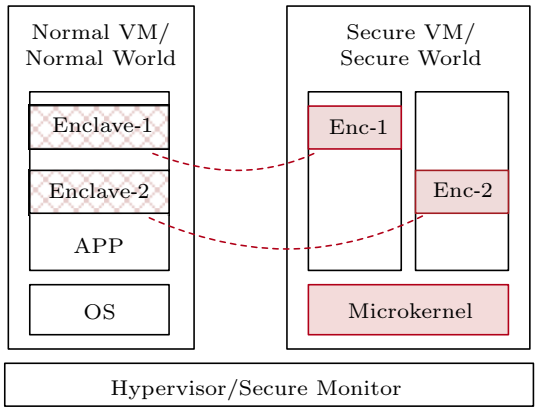

Fig. 3. UniTEE gives SGX-like enclave abstraction based on AMD SEV or ARM TrustZone. Enclave is abbreviated as Enc.

An application still runs on the untrusted OS while its enclaves are created by and run on the microkernel. The microkernel is only responsible for the enclave life-cycle management, which mainly involves enclave construction/destruction, enclave memory management, and enclave thread scheduling. The application cannot access its enclaves' memory while an enclave can access its host application's memory only if the microkernel maps the normal memory belonging to the application into the enclave's address space. By default, an enclave and its host application have shared memory for communication.

In brief, UniTEE leverages the microkernel to multiplex a single hardware TEE and thus allows an application to create an arbitrary number of fine-grained enclaves on an AMD SEV machine or an ARM TrustZone machine, just like on an Intel SGX machine. The tradeoff is that our microkernel enlarges the trusted computing base (TCB) of an enclave. Nevertheless, our microkernel has a small code base (around 5 000 lines of code) and thus is relatively easier to be implemented correctly. With more efforts in the future, formal verification can be used to make our microkernel more secure. Although prior work on ARM TrustZone also proposed to deploy a secure OS in the secure world (e.g., OPTEE), the secure OS is for running multiple trusted applications instead of enclaves (belonging to the host applications), which makes the secure OS and the microkernel of UniTEE different.

3.3 Programming Interfaces

As the SGX programming model is easy to use and adopted by the public, UniTEE preserves similar (or even can be the same) interfaces as listed in Table 1. By providing such interfaces, existing secure applications targeting SGX can be more easily ported to UniTEE, which can make our work more practical.

3.3.1 Enclave Creation

Fig.4 shows how a host application creates an enclave on SEV and TrustZone platforms (the enclave creation procedure on SGX platforms is just like before, i.e., using official Intel SGX Driver). First, the host application prepares the enclave image and the corresponding configuration. Second, it invokes create_enclave, which traps into a kernel module (Drv-1) deployed by UniTEE. Third, the kernel module transfers the control flow to the microkernel. For transferring the control flow, the SMC instruction is used on TrustZone-enabled platforms while VMMCALL and VMRUN instructions are used on SEV-enabled platforms. In the former case, the SMC instruction makes the CPU trap into the monitor mode, and another tiny module (Drv-2) deployed by UniTEE helps to finish the switch between the normal world and the secure world. In the latter case, the VMMCALL instruction triggers a VMExit and makes the CPU trap into the hypervisor mode and, thus, a similar tiny module (Drv-2) in the hypervisor executes the VMRUN instruction to notify

Table 1. Main Interfaces in Enclave-Management Library (for Host Applications) and Modified C Library (for Enclaves)

Declaration	Description
int create_enclave(Buf enclave_img, Buf enclave_config)	Used by the host application to create an enclave. The two arguments give the locations of the enclave image and the configuration, respectively.
int attest_enclave(int enclave_id, Buf input, Buf output)	Used by the host application to generate an attestation for an enclave. The last two arguments give the locations of the input message and the final attestation data, respectively.
int call_enclave(int enclave_id, Buf buffer) [ecall]	Used by the host application to invoke an enclave function. The first argument specifies which enclave to call. The second one is the shared buffer between the host application and the enclave, which is used for storing both input arguments and output results.
int call_host(Buf buffer) [ocall]	Used by an enclave to invoke its host application's function.
int get_seal_key(Buf output_key)	Used by an enclave to get a sealing key which can be used to encrypt some persistent data outside the enclave.
Most interfaces in musl-libc	An enclave can also invoke common POSIX interfaces in musl-libc just like a normal application.

the microkernel. Fourth, a system service of the microkernel, named Enclave Construction Service (ECS), receives the enclave creation request and then constructs the enclave according to the image and configuration passed through the shared memory. On SGX platforms, UniTEE provides the same interface but constructs enclaves just like how official Intel-SGX SDK does. Specifically, a kernel module like Drv-1 in the OS builds enclaves using SGX instructions (ENCLS).

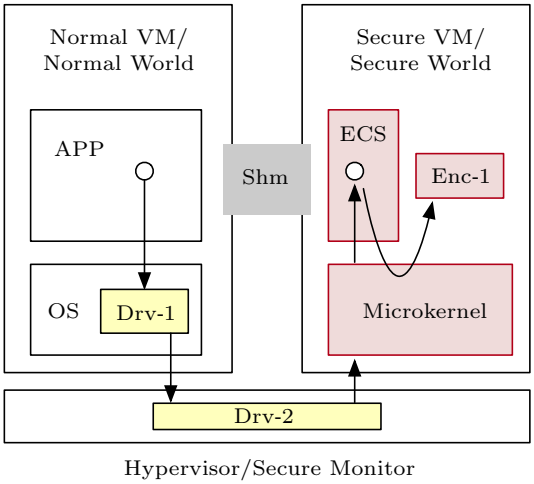

Fig.4. Procedure of enclave construction on SEV and TrustZone platforms. ECS represents enclave construction service. Enclave is abbreviated as Enc. Drv-1 and Drv-2 are two software components deployed by UniTEE.

Components of UniTEE. For SEV, the components include a kernel module in the normal VM's guest OS, a tiny module in the hypervisor, and a microkernel OS in the secure VM. For TrustZone, the components include a kernel module in the normal world OS, a tiny module in the monitor mode, and a microkernel OS in the secure world. For SGX, the components include a kernel module, i.e., the SGX driver. Besides, the components also include a library in each enclave for all the three platforms.

3.3.2 *Enclave Attestation*

Remote attestation enables a remote user to attest whether an enclave is correctly launched and further allows the remote user and the enclave to build a secure communication channel (i.e., exchanging a session key). UniTEE provides the corresponding interface named *attest_enclave*. On SGX-enabled platforms, UniTEE simply uses the hardware-provided remote attestation mechanism. On the other two platforms, UniTEE leverages the reliable microkernel to implement a two-phase attestation. Briefly speaking, in the first phase, it boots the microkernel by using the secure boot mechanism provided by the hardware and allows remote users to negotiate secure keys with the microkernel; in the second phase, the ECS of the microkernel is responsible for launching enclaves, computing the enclave measurement, and signing the measurement by secure keys. As the results of *attest_enclave*, the signed measurement will be returned to the host application and it can be further sent back to remote users for attestation.

3.3.3 *Enclave Interaction*

An enclave and its host application can expose function routines for each other, called *ecall* and *ocall* in the official Intel SGX SDK. UniTEE also supports the interaction between enclaves and the host application, and the programming interfaces are *call_enclave* and *call_host*. When the host application thread invokes an enclave function, it needs to transfer the control flow to an enclave thread. As shown in Fig.5(a), UniTEE can implement the cross-boundary invocation (passive

interaction mode) by using a similar method for creating an enclave. However, in such mode, the invocation cost is high because both the SMC-based world switching of TrustZone and the VMMCALL/VMRUN-based VM switching of SEV bring both expensive direct cost (thousands of CPU cycles) and indirect cost (pollution to CPU internal structures like cache and TLB). For the sake of performance, UniTEE also provides an alternative interaction mode (proactive interaction mode) which integrates the FlexSC-like mechanism [49]. There is a shared buffer between a host application thread and an enclave thread. The buffer is not only for transferring the data but also for transferring the control flow. The enclave thread will poll on a request ready flag in the buffer. When the flag is set, which means the application thread makes an invocation request, the enclave thread starts to handle the request and sets a replay ready flag after finishing the request. Therefore, when invoking *call_enclave*, the host application thread writes the arguments of the request to the shared buffer, sets the corresponding request ready flag, and waits for the reply ready flag. After the enclave thread sets the reply ready flag, it retrieves the results of the request and continues the execution. During the request handling procedure, the enclave thread may also invoke functions provided by the host application through *call_host*. If so, it sets the reply ready flag to a specific value, which means it makes an invocation request to the host application thread. Since the latter thread polls on the reply ready flag, it can detect and finish such a request. With this proactive interaction mode, the invocation can be much faster while it requires more CPU cores. The application programmers can select either mode according to the requirements.

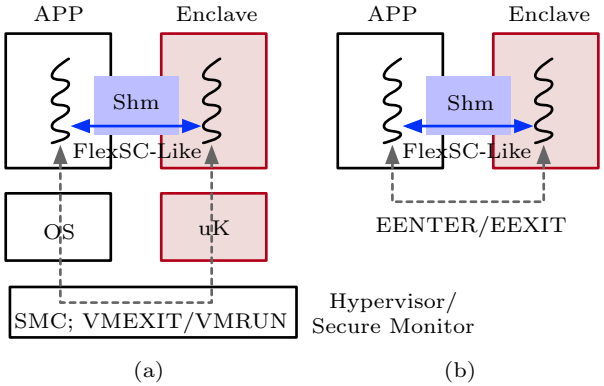

Fig.5. UniTEE supports two modes of communication between an enclave and its host application. The shared communication buffer is abbreviated as Shm.

Fig.5(b) shows that UniTEE also enables the two modes of interaction on SGX platforms. In the passive interaction mode, expensive *EENTER* and *EEXIT* instructions are used. In the proactive interaction mode, the FlexSC-like mechanism described above is used instead.

3.3.4 *Enclave Library*

UniTEE provides an in-enclave C library based on musl-libc for easing programming and supporting running legacy code inside an enclave. The vanilla musl-libc will finally invoke system calls by executing *syscall* (x86-64) or *svc* (AArch64) instruction. Unlike that, the modified library changes the system calls into invocations to the host application and then the host application will invoke the requested system calls on the OS for the enclave. In other words, the system calls issued by an enclave are redirected to the OS on which the host application runs. Note that an enclave belongs to its host application, and the OS will serve it for most system calls. Although the microkernel of UniTEE does implement various system calls like an OS, it provides the ones related to enclave memory management. The enclave library will transparently dispatch the system calls without the involvement of programmers. An existing application can be linked against the modified C library and then directly run in an enclave as a whole. In this case, UniTEE will start a simple host application which just creates the enclave and handles the system calls at runtime for the enclave.

Besides, the library supports another interface named *get_seal_key* which can be invoked by an enclave to get a sealing key. If an enclave needs to store some data for use in the next boot, it can seal the data with this key and store the encrypted data on some untrusted storage. Enclave Key Service, another system service of the microkernel, manages the relationship between the sealing key and the enclave measurement. Therefore, the same enclave (i.e., the same measurement) can retrieve the same sealing key after every boot. For SGX platforms that directly support this functionality, the enclave library uses the *EGETKEY* instruction to get the sealing key.

4 Heterogeneous TEE Migration

The unified programming abstraction of UniTEE benefits secure applications development by hiding the heterogeneity from programmers. Furthermore, based on the unified abstraction, UniTEE transparently enables the enclave migration between different platforms.

We focus on the enclave migration in this paper because: the migration of an application's non-enclave part has no significant difference from traditional migrations, which has been detailedly presented in prior studies [22, 23]. In this section, we first give an overview of the whole enclave migration process and then explain the detailed techniques used during the migration.

4.1 Overview of Migration

Briefly speaking, as shown in Fig. 6, a migration process includes the following three steps: first, an enclave checkpoint is generated on the source machine; second, the checkpoint is transferred to the target machine through the network; third, the checkpoint is used to restore the running states and resume the execution of the enclave on the target machine.

Compared with the traditional checkpoint generation of application migration, there are three differences that make the enclave checkpoint generation challenging. 1) The enclave states cannot be accessed by any system software (e.g., OS), which means they cannot help to generate the enclave checkpoint; 2) the system software may be compromised and launch consistency attacks during the generation procedure; 3) the instructions and calling convention used by enclaves are different on heterogeneous TEEs.

To overcome the first challenge, UniTEE enables each enclave to generate its own checkpoint, i.e., an enclave encrypts and then dumps out its states as a checkpoint without the involvement of others like the OS. Considering state consistency, an enclave needs to stop the enclave threads before generating the checkpoint. Otherwise, the checkpoint may be inconsistent, i.e., it consists of both old and new data. Also, a malicious OS may schedule enclave threads during checkpoint generation to break the state consistency. To overcome this second challenge, UniTEE enables an enclave to make all its threads enter into a quiescent point (make no further updates) before checkpoint generation. Besides, the underlying hardware TEEs on the source and the target machines can be heterogeneous. UniTEE integrates the heterogeneous migration techniques proposed in Popcorn [23] to solve the heterogeneity challenge (the third one). During the generation process, UniTEE will transform the architecture-dependent states according to the target machine's architecture.

Before receiving the enclave checkpoint, the target machine will first launch a virgin enclave with the enclave binary for its architecture. The virgin enclave will receive the checkpoint from the source enclave and use the checkpoint to resume the execution. For securely transferring the checkpoint from the source enclave to the target enclave, the two enclaves will negotiate a migration key with each other by using the widely-used Diffie-Hellman key exchange protocol whose crux is the mutual authentication between the two participants. As UniTEE enables remote attestation (introduced in Subsection 3.3), the source enclave and the target enclave can attest each other to finish the key exchange protocol and then generate the migration key (stored inside the enclaves). Before writing the checkpoint out, the source enclave will first calculate the checksum of the checkpoint and then encrypt it together with the checksum by using the migration key. Since the checkpoint is encrypted when it is outside the enclave or in the network, the untrusted software like OS cannot break confidentiality and integrity.

4.2 Preparation for Checkpointing

UniTEE introduces control thread, an extra enclave thread, to assist migration. Since the control thread runs within an enclave, it can traverse and dump the entire memory data within the enclave boundary as the

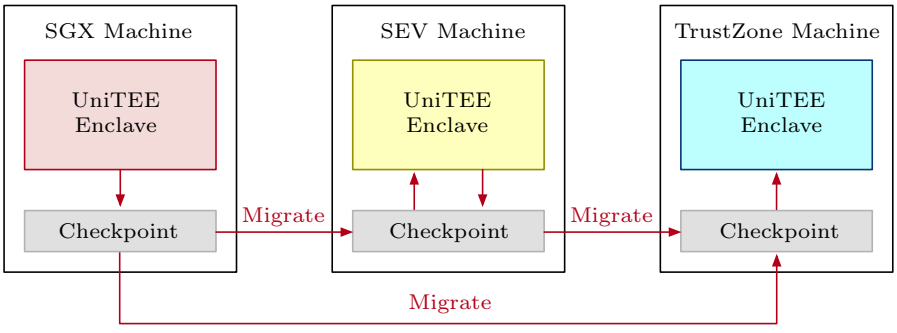

Fig.6. Secure enclave migration between different platforms.

checkpoint. To ensure state consistency of the generated checkpoint, it has to make all the other enclave threads (worker threads) suspend running before starting the generation. Otherwise, it may get a checkpoint with inconsistent data because a worker thread may update some memory during the generation process. As a user-level thread, the control thread cannot directly suspend all worker threads. However, if it asks the OS for help, a malicious OS can deceive the control thread that all enclave threads are suspended but actually not, which will violate the consistency of checkpoint.

Fig. 7 presents a simple example of such a data consistency attack. When a migration begins, a worker thread in an enclave is transferring money from account A to account B. The control thread calls *stop_other_thread*() to ask the OS to stop all other enclave threads. However, the malicious OS returns OK but actually does not stop the worker threads. Thus, the control thread may get an old version of account A (5 000) and a new version of account B (5 000), which violates the invariant that the sum of accounts should be 5 000.

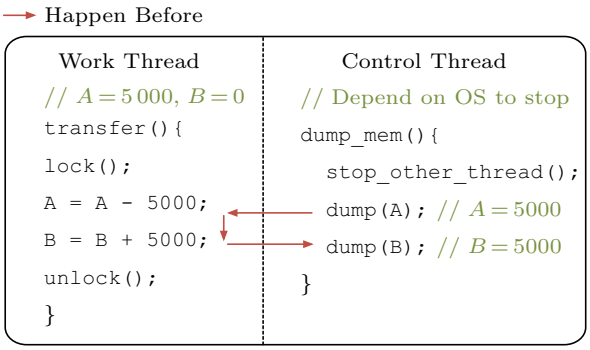

Fig.7. Example of data consistency attack.

Instead of relying on the untrusted OS, the control thread makes the worker threads reach a quiescent point as follows. When receiving a migration notification (e.g., through a user-defined signal like SIGUSR1), the control thread is wakened up and then sets a global flag in the enclave to indicate the start of the suspending process. There is a global flag for each enclave and a local flag for each worker thread. Initially, the global flag is unset, and the local flags are "free". Each worker thread sets its local flag to "busy" and "free" at the enclave entry point and the exit point, respectively. Each worker thread normally runs until meeting a migration stub (see Fig. 8), in which it first checks whether the global flag is set. If not, it continues to execute normally. If so, it performs stack transformation, sets its corresponding local flag to "spin", and then enters the spin region. The stack transformation is for transforming the execution stack according to the target architecture, which will be explained in Subsection 4.3. When running in the spin region, a worker thread will not change any memory and will keep in the region, until it finds that the global flag is unset. The control thread will wait for the point when all the local flags of worker threads are either "free" or "spin" (i.e., not running or in the spin region) before generating the enclave checkpoint. Therefore, it can ensure the consistency of the checkpoint without the help of the untrusted OS.

```
void migration_stub(void)
{
  if (global_flag == set) {
        transform_stack();
        local_flag = spin;
        while (global_flag == set);
  }
}
```

Fig.8. Pseudo-code of the migration stub.

UniTEE inserts migration stubs before the ocalls. Therefore, a running enclave thread will respond to the migration when it invokes an ocall. Nevertheless, some worker threads may execute in the enclave for a long time without performing an ocall. It is very likely that such a thread has already set its local flag to "busy" when the control thread sets the global flag. If so, the control thread needs to wait for a long time, which will block the process of migration. To this end, UniTEE also allows programmers to insert migration stubs in their code as they want.

4.3 Hiding Heterogeneity for Migration

Since the three popular hardware TEEs, namely, Intel SGX, AMD SEV, and ARM TrustZone, are provided on different architectures, UniTEE also has to transparently hide the heterogeneity during the enclave migration. We explain the detailed techniques from the following four main aspects.

1) *How to Migrate Code.* Different hardware TEEs must use the corresponding CPU instructions. Therefore, UniTEE compiles different enclave binary codes for different hardware TEEs, and an enclave will use the corresponding binary code according to the underlying hardware TEEs. The key point of the compilation is that each function in the different binary is located at the same start address. Therefore, function point-

ers are always valid after migration, which eases the migration, i.e., no need to update the pointers.

2) *How to Migrate Data.* UniTEE targets 64-bit and little-endian because the first two types of TEEs support 64-bit only and all three types use little-endian. Thus, the data format needs no transformation across the three architectures, e.g., the data format of a struct written in C is always the same for the three TEEs. Like the function start address, each global variable address is also located at the same address. Therefore, the global data section and the heap area (using the same heap start address) can be directly copied from the source enclave to the target enclave. The validity of data pointers is inherently preserved after migration, which significantly eases the migration process.

As shown in Fig.9, UniTEE generates multiple enclave binaries for each architecture. An enclave binary can be distributed to the target machine before migration or when migration is triggered. When a migration begins, the target machine boots the virgin enclave that will receive the checkpoint from its source enclave. Thus, the code section does not need to be transferred during migration. As described above, the enclave binaries for different TEEs share a uniform address space layout, i.e., every symbol address is kept the same, and every data structure uses the same memory format. Therefore, simply copying the data section and heap area will not make any pointer invalid. In other words, these two areas are transferred without any transformation.

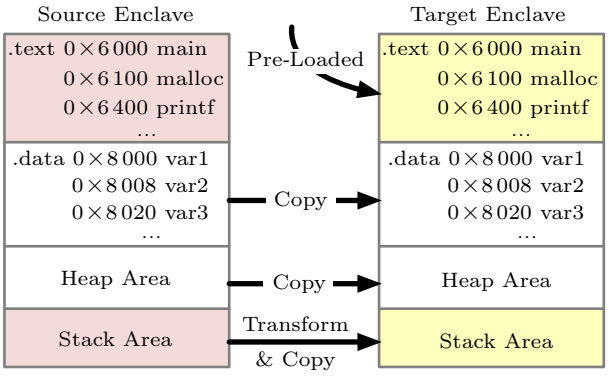

Fig.9. Memory layout of the source enclave and the target enclave. Colored parts are architecture-dependent.

3) *How to Migrate Execution Context as Well as Execution Stack.* Different architectures provide different numbers of general-purpose registers (GPR) and use different calling conventions. For example, there are 16 GPRs for Intel SGX (x86-64) but 32 GPRs in ARM TrustZone (AArch64). And the calling conventions for them vary widely, which makes the execution stack different. A simple approach to solving this challenge is to use the same number of registers and the same calling convention. In such a way, it is easy to give an explicit one-to-one mapping to connect registers in different ISAs and simply copying the stack area can also work for migration. However, this approach leaves many registers unused and abandons originally applicable optimizations, which may hurt the performance enormously.

Instead, UniTEE adopts and implements the stack-/register transformation proposed by prior work [23]. The basic idea is recording each stack frame's information at compile time and then reconstructing the execution stack frame by frame for the target architecture at migration time. The inserted migration stubs ensure a migration always happens at the function boundaries, which means the stack frame to transform is always intact. Specifically, the compilation toolchain is based on LLVM and the information (including live variables and the calling site) of each stack frame is recorded according to the intermediate representation (IR) of LLVM. For each specific architecture, a live variable is either mapped to a register or on the stack. Therefore, according to such information, the transformation procedure will reconstruct a new stack and set the registers for the target architecture.

Implementation Details for the Transformation. During the compilation process, all the stack frame information is recorded in a particular section of the binary, which mainly includes locations of live variables (either on the stack or in a register). The transformation procedure first calculates the size for the new stack according to the recorded information. Then it rebuilds the new stack from the outermost frame to the innermost frame (frame by frame). For rewriting one stack frame, it gets all the live variables of that frame in the source enclave binary, queries their locations in the target enclave binary, and then copies them to the new locations. A special case is that a variable is a pointer that points to some stack address. In this case, the variable cannot be directly copied because its value (some stack address) should be changed after the stack is transformed. Instead, it will be recorded in a fixup list for resolving later. Every time when the stack transformation procedure copies a variable, it checks whether the variable ($V1$) is pointed by some variable ($V2$) in the fixup list. If so, it removes $V2$ from the list and sets the new location of $V1$ (on the target enclave

stack) to *V2* for the target enclave. Also, the return address of each stack frame is also rewritten according to the calling site information.

At the migrate point, some live variables may be stored in registers. Therefore, besides transforming each stack frame, it is also necessary to restore these in-register variables for the target enclave. According to the information recorded during compilation, the transformation procedure knows the location of each live variable on different architectures and thus can simply set the corresponding registers for the target enclave.

4) *How to Migrate OS-Related States.* UniTEE also allows an enclave to invoke system calls, as described in Subsection 3.3. Therefore, it is also necessary to migrate OS-related states from the source enclave to the target enclave. Currently, UniTEE supports restoring file descriptors and TCP connections. For file descriptors, it records the states (e.g., file path, file descriptor, and cursor) of each opened file in the modified enclave library that redirects the system calls. These states and related files are also transferred during migration. During the restoring process, the target enclave reopens each file and sets the cursor to the right position. For migrating TCP connections, UniTEE refers to CRIU[③]. The Linux kernel (since version 3.5) has supported the TCP connection repair mechanism to help with sockets transmission. UniTEE first uses TCP_REPAIR option to switch the socket into a special mode for the source enclave. It then collects and transfers necessary TCP states. Last, on the target side, the TCP states will be restored, and the socket mode will be reset to normal for the target enclave.

5 Evaluation

We conduct performance evaluations on the prototype of UniTEE and present the results in this section. The experiments include the enclave migration between Intel SGX and ARM TrustZone, between Intel SGX and Intel SGX, and between Intel SGX and AMD SEV.

5.1 Between SGX and TrustZone

We conduct experiments on two machines that support Intel SGX and ARM TrustZone, respectively. One is equipped with Intel Core i7-9700 CPU and 16 GB memory, and the other is an HiKey970 board with 6 GB memory. We run Ubuntu 16.04 on both machines while the Linux kernel versions are 4.15.0 and 4.9.78. We run each experiment over 30 times and report the average of the results. The standard deviation is within 5% across all the experiments.

We select several SPEC CPU 2006 benchmarks and vedis (a popular key-value store) as applications. An application is protected as a whole and runs in an enclave of UniTEE. We do not modify the source code of the applications except for inserting some migration stubs. For the SPEC CPU benchmarks, the workloads are the built-in ref test suites. For vedis, we generate 10 million random keys and perform PUT and GET operations (50% PUT and 50% GET) randomly.

Overhead of Migration Support. We first present an experiment on the overhead introduced by supporting migration. We compile the chosen benchmarks with and without migration support, run them in enclaves, and measure the execution time. To disable migration support, we do not link the applications with migrated-related libraries or insert any migration stubs. Fig.10 shows the results of some benchmarks (others are similar). We normalize the results to the no-migration-support version for better readability.

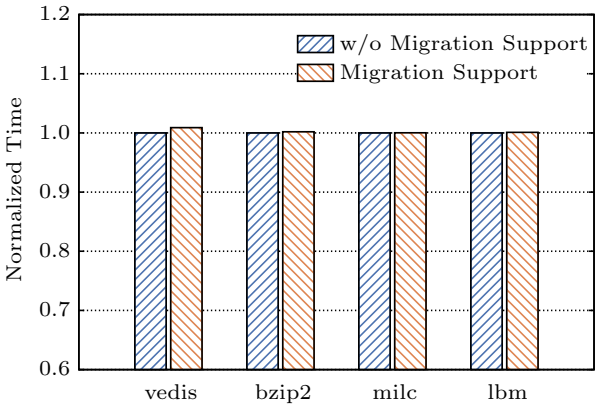

Fig.10. Execution time with and without (w/o) migration support. The results are normalized to the no-migration-support version.

Migration Cost. Table 2 gives the breakdown time of migrating applications from the Intel machine to the ARM machine. To facilitate analysis, we divide the whole migration procedure in both machines into different phases: preparation phase, checkpoint phase, transmission phase, and restore phase. The time consumed by booting the virgin target enclave is not reported because this procedure is out of the critical path, i.e., having finished before the final restore phase.

This evaluation shows the migration support brings nearly zero overhead. It makes sense because UniTEE

[③]CRIU: Checkpoint/Restore In Userspace. https://www.criu.org/Main_Page, Feb. 2021.

only requires each thread to do two extra operations during normal execution for supporting migration. The first one is to initialize the local migration flag, which can be ignored since it only happens once. The second one is to check the global migration flag at each migration point. However, checking the flag only requires several instructions, which is negligible compared with real workloads. Therefore, migration support does not influence performance during normal execution.

Table 2. Breakdown of the Migration Cost

Benchmark	Preparation Phase (µs)	Checkpoint Phase (µs)	Transmission Phase (µs)	Restore Phase (µs)
vedis	932	155 818	1 140 314	71 888
bzip2	543	618 690	2 526 885	162 563
milc	916	399 044	1 880 252	113 955
sjeng	713	529 231	2 173 368	150 601
libquantum	755	294 514	1 378 012	90 352
h264ref	598	206 130	1 261 305	76 975
lbm	701	1 251 807	4 629 325	410 445

Note: Preparation and checkpoint phases happen on the source machine. The transmission phase is for transferring the enclave checkpoint through the network. The restore phase happens on the target machine.

Preparation Phase. This phase mainly consists of the time of waiting for the quiescent point and performing stack/register transformation. As shown in Table 2, it takes less than 1% of total migration time in all benchmarks. Here each enclave only has one worker thread, thereby the quiescent point is reached once it meets the first migration point. Stack transformation can be done in a short time because all necessary information has been stored in binaries during compilation and the depth of the stack is usually not deep.

Checkpoint/Restore Phase. After the preparation phase, the enclave (control thread) encrypts its memory data and dumps the encrypted data outside the enclave, which is called the checkpoint phase here. Similarly, for restoring the checkpoint in the target enclave, the checkpoint needs to be copied into the enclave and then decrypted. For better performance, the checkpoint only contains the enclave memory in use. Firstly, the code section of the enclave does not need to be dumped, since the binary code for different architectures is different and can be placed on the target machine in advance. Secondly, the size of the data section is known at compile time and does not change at runtime. The data section is included in the checkpoint. Thirdly, only the in-use enclave heap region is dumped to the checkpoint. To be specific, the in-use heap region consists of two parts: one is from the heap base to the heap top; the other is a list of memory-mapped areas (i.e., through mmap). Fourthly, only the valid stack regions are dumped according to the stack pointers.

The cost of the checkpoint phase (on the Intel SGX machine) is obviously higher than that of the restore phase (on the ARM TrustZone machine). This is because accessing the SGX-protected memory is much more expensive, especially when SGX page swapping happens.

The cost of such two phases is related to the in-use enclave memory size. Therefore, we give a further analysis of the cost and Fig.11 shows the results. Fig.11(a) and Fig.11(b) show the time spent on encryption and decryption. The encryption mechanism is advanced encryption standard (AES) and some hardware-assisted acceleration can be used. For example, Fig.11(a) shows Intel AES-NI [50] instructions can reduce over 60% of encryption time. The time consumed in dumping (writing the checkpoint to outside) and restoring (copying the checkpoint into an enclave) is reported in Fig.11(c) and Fig.11(d), respectively. The result shows the dumping and restoring time grows linearly as the enclave size increases, because dumping and restoring are actually memory copy operations.

Transmission Phase. Fig. 12 presents the time of transferring enclave checkpoints with different sizes. The two machines connect to the same LAN, and the bandwidth is about 115 MB/s. The evaluation results show the time consumed in the transmission phase increases as the enclave size grows. According to Table 2, this phase takes most of the migration time (over 78%). With a faster network, the total migration latency can be significantly reduced.

Overlapping Phases. To further decrease the migration latency, we pipeline the execution of the checkpoint/transmission/restore phases by dividing the whole enclave checkpoint into pieces. Fig.13 shows the time saved by using this strategy on the source machine. The total latency can be reduced by up to 30%. Therefore, the total downtime for the enclave applications is bottlenecked/decided by the network speed. We conclude the solution proposed by UniTEE is feasible.

5.2 Between Two SGX Machines

We also measure the performance of enclave migration between two laptops with Intel Core i7-6700HQ 2.6 GHz CPU and 8 GB memory. The experiment is

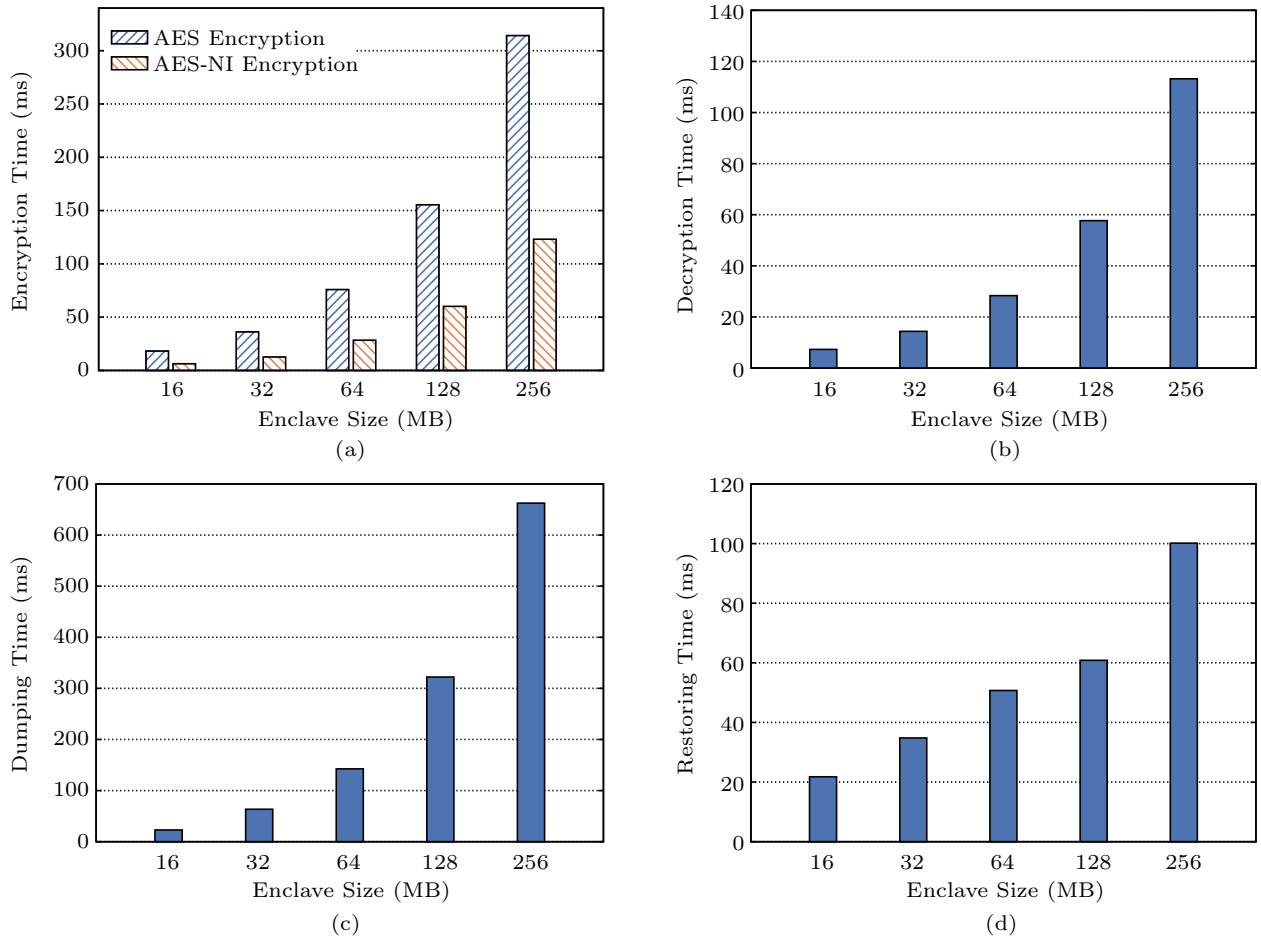

Fig.11. Time consumed in the checkpoint/restore phase. (a) Encryption. (b) Decryption. (c) Dumping. (d) Restoring.

migrating a virtual machine (VM) with and without enclaves running inside. KVM is chosen as the underlying hypervisor, and the version of QEMU is 2.5.0. The guest VM has four VCPUs (virtual CPUs) and 2 GB memory.

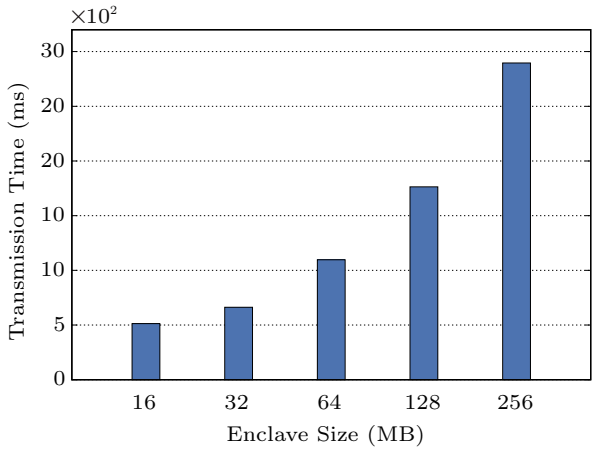

Fig.12. Network transmission time.

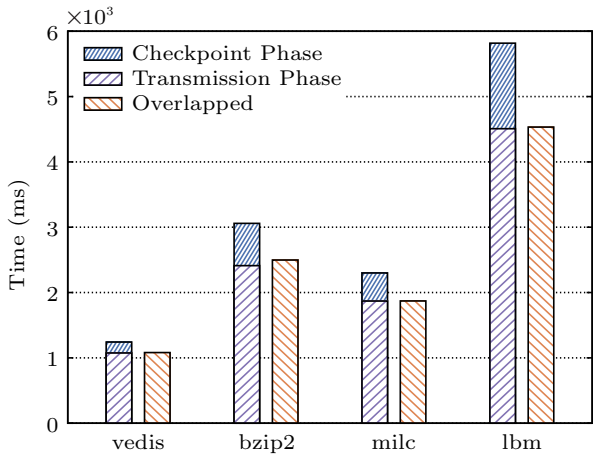

Fig.13. Time saved by overlapping phases in the source machine.

We run two VMs respectively, one with some running enclave applications, and the other with the same number of original applications. The enclaves run either *libjpeg* or *mcrypt*, which are real-world applications. The enclave size is 1 MB, and the workload is an

endless loop of picture decoding or encryption. Fig.14 shows the total migration time. The migration of VM with no more than 32 enclaves has about 2% overhead. The overhead increases to 5% when the number of enclaves reaches 64.

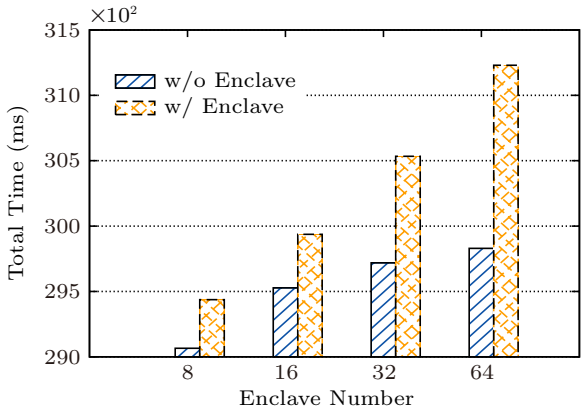

Fig.14. Total migration time with (w/) and without (w/o) enclaves. Note that the x-axis does not start from 0.

Besides, we conduct an experiment to compare the SDK performance, i.e., UniTEE and Intel official SGX SDK. The benchmark is nbench 2.2.3 in which most applications are computation-intensive. String-sort is the one that accesses much more secure memory, which leads to high SGX paging overhead. As shown in Fig.15, our SDK shows similar or better performance compared with Intel SDK.

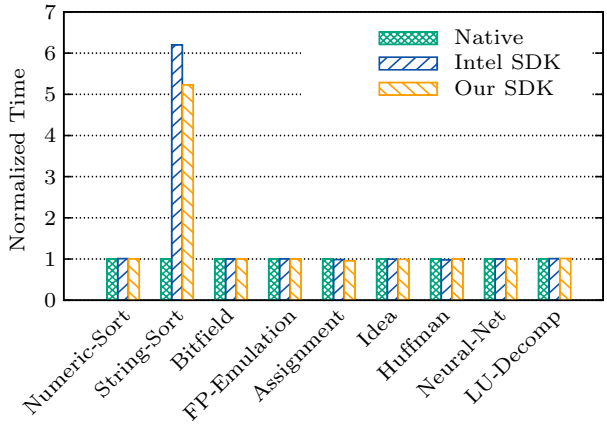

Fig.15. SDK performance comparison on nbench.

5.3 Between SGX and SEV

Intel SGX and AMD SEV are two security extensions to x86-64 and they share the same general purpose registers as well as calling convention. Compared with the migration between SGX and TrustZone, the migration between SGX and SEV needs no stack transformation while the other procedures are the same. Fig.16 shows the time for generating the enclave checkpoint. The AMD machine is equipped with EPYC 7281 CPU that supports SEV.

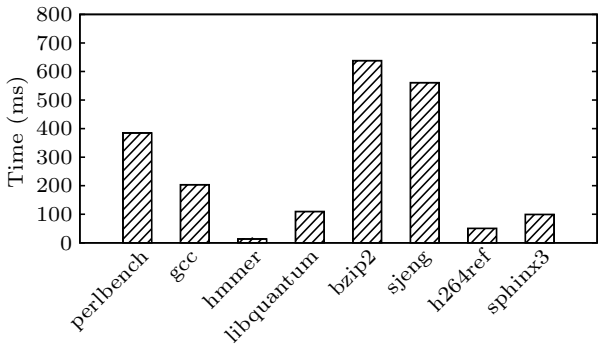

Fig.16. Cost for generating checkpoints.

Besides, we measure the execution time of 12 applications in SPEC CPU 2006 benchmarks on the SEV machine and present the results in Fig.17. This experiment is to show the protection overhead of UniTEE on the SEV machine. For CPU-intensive benchmarks such as bzip2 and gobmk, the performance of enclaves is nearly the same as or even better than the native execution performance. Two reasons can explain: first, UniTEE will not bring overhead to enclave applications when they do not invoke ocalls; second, when executing ocalls (system calls), the FlexSC-like design (described in Subsection 3.3) avoids the context switches for the enclave threads although it requires some extra cycles. Context switches between user-mode (ring 3) and kernel-mode (ring 0) may incur indirect costs like TLB/cache pollution. Other applications show less than 5% overhead which mainly comes from the memory copies (transferring arguments and results) during ocalls.

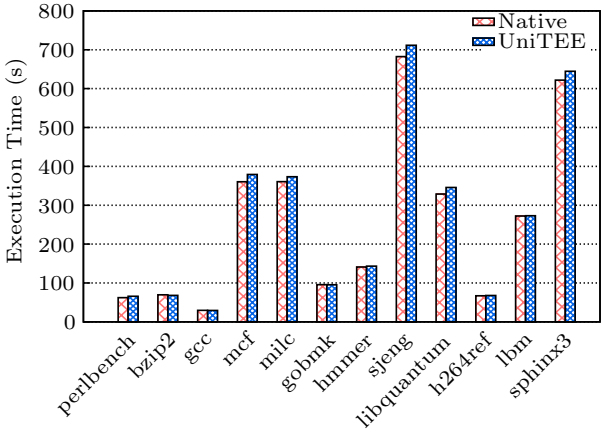

Fig.17. Runtime overhead of UniTEE on the SEV machine.

6 Discussion

As mentioned in Subsection 3.2, the microkernel used in UniTEE is included in the TCB. It makes sense to assume the microkernel is trusted because it only provides simple and clear functionalities and has a small code base. This is also a common assumption (e.g., TrustVisor [51], CloudVisor [52], Nested Kernel [53]). It may also be feasible to build UniTEE's microkernel over a formally verified OS kernel (e.g., Hyperkernel [54] and seL4 [55]) and further verify the entire one with more effort. Besides, the implementations of existing hardware enclaves also heavily rely on software. For example, as mentioned in "Hardware is the new software" [56], much of SGX's logic is implemented by microcode, which can be patched on-the-fly just as software. Meanwhile, many researchers try to build enclaves from software (with the help of hardware), like Komodo [57]. We certainly agree that from the perspective of security, it is more preferable to construct the TCB in a more simple and predictable way. But we argue that the point here is more about the level of semantics instead of being hardware or software. In our design, we try to move some of the hardware logic from firmware (e.g., on AMD PSP) to software running in the secure TEE, which has low-level semantics, instead of developing a new complex software like a guest OS.

7 Related Work

Enclave Programming Model. The strong security insurance of hardware TEEs motivates a variety of prior studies [18, 19, 21, 35, 44, 58, 59] to protect applications by leveraging one of Intel SGX, ARM TrustZone, and AMD SEV. However, they do not focus on providing a unified enclave programming model for hiding the underlying hardware security technologies. Open Enclave SDK④ aims to provide consistent API surface across enclave technologies as well as all platforms from cloud to edge, which shares the same goal of the unified programming abstraction of UniTEE. Nevertheless, Open Enclave SDK considers SGX and TrustZone while UniTEE further considers SEV. Moreover, UniTEE enables enclave migration across those platforms while Open Enclave SDK does not.

Microkernel Usages. There is a long line of research on microkernel OS [55, 60–64]. Owing to the desired advantages including good security and fault isolation, microkernel has been used in some safety-critical scenarios like vehicles. Designing microkernels for general-purpose scenarios is also on the way. Nevertheless, UniTEE leverages the microkernel for constructing isolated enclave instances in a single hardware TEE. A recent work [65] also proposes to design a TEE OS based on the microkernel architecture. Different from that, the microkernel of UniTEE only manages the enclave life cycle without providing various OS services through system calls because most system calls are redirected to the full-fledged OS which runs the host application. The microkernel used here is derived from [63].

Heterogeneous Migration. Live migration between heterogeneous architectures has been studied by prior work [11, 22, 66, 67]. UniTEE adopts and extends the existing migration techniques of Popcorn [11, 22] to migrate the secure enclaves among different TEE hardware. The major difference is that the OS is not trustworthy. During migration, UniTEE relies on the OS functionalities without trusting it. For example, UniTEE requires the OS to transfer the enclave checkpoint through the network but protects the consistency, confidentiality, and integrity of the checkpoint. [24] makes efforts to securely migrate SGX enclaves on untrusted cloud, which, however, does not provide unified enclave abstraction or enclave migration support on heterogeneous platforms.

8 Conclusions

This paper proposed UniTEE whose target is twofold. It provides a unified enclave programming abstraction that can help programmers to write enclave applications without considering the underlying hardware TEEs. Further, with the unified abstraction, it enables secure enclave migration between heterogeneous platforms.

UniTEE could be extended to more heterogeneous security architectures like confidential VMs, which is our future work.

References

[1] Park H, Zhai S, Lu L, Lin F X. StreamBox-TZ: Secure stream analytics at the edge with TrustZone. In *Proc. the 2019 USENIX Annual Technical Conference*, July 2019, pp.537-554.

[2] Shi W, Cao J, Zhang Q, Li Y, Xu L. Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 2016, 3(5): 637-646. DOI: 10.1109/JIOT.2016.2579198.

[3] Hu Y C, Patel M, Sabella D, Sprecher N, Young V. Mobile edge computing—A key technology towards 5G. Technical

④Open Enclave SDK. https://openenclave.io/sdk/, Feb. 2021.

Report, European Telecommunications Standards Institute, 2015. https://infotech.report/Resources/Whitepapers/f205849d-0109-4de3-8c47-be52f4e4fb27_etsi_wp11_mec_a_key_technology_towards_5g.pdf, Dec. 2021.

[4] Satyanarayanan M. The emergence of edge computing. *Computer*, 2017, 50(1): 30-39. DOI: 10.1109/MC.2017.9.

[5] Shi W, Dustdar S. The promise of edge computing. *Computer*, 2016, 49(5): 78-81. DOI: 10.1109/MC.2016.145.

[6] Stojkoska B L R, Trivodaliev K V. A review of Internet of Things for smart home: Challenges and solutions. *Journal of Cleaner Production*, 2017, 140: 1454-1464. DOI: 10.1016/j.jclepro.2016.10.006.

[7] Nastic S, Rausch T, Scekic O, Dustdar S, Gusev M, Koteska B, Kostoska M, Jakimovski B, Ristov S, Prodan R. A serverless real-time data analytics platform for edge computing. *IEEE Internet Computing*, 2017, 21(4): 64-71. DOI: 10.1109/MIC.2017.2911430.

[8] Machen A, Wang S, Leung K K, Ko B J, Salonidis T. Live service migration in mobile edge clouds. *IEEE Wireless Communications*, 2017, 25(1): 140-147. DOI: 10.1109/MWC.2017.1700011.

[9] Wang S, Xu J, Zhang N, Liu Y. A survey on service migration in mobile edge computing. *IEEE Access*, 2018, 6: 23511-23528. DOI: 10.1109/ACCESS.2018.2828102.

[10] Islam M, Razzaque A, Islam J. A genetic algorithm for virtual machine migration in heterogeneous mobile cloud computing. In *Proc. the 2016 International Conference on Networking Systems and Security*, Jan. 2016. DOI: 10.1109/NSysS.2016.7400696.

[11] Barbalace A, Karaoui M L, Wang W, Xing T, Olivier P, Ravindran B. Edge computing: The case for heterogeneous-ISA container migration. In *Proc. the 16th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments*, Mar. 2020, pp.73-87. DOI: 10.1145/3381052.3381321.

[12] Rodrigues T G, Suto K, Nishiyama H, Kato N, Temma K. Cloudlets activation scheme for scalable mobile edge computing with transmission power control and virtual machine migration. *IEEE Transactions on Computers*, 2018, 67(9): 1287-1300. DOI: 10.1109/TC.2018.2818144.

[13] Roman R, Lopez J, Mambo M. Mobile edge computing, fog *et al.*: A survey and analysis of security threats and challenges. *Future Generation Computer Systems*, 2018, 78: 680-698. DOI: 10.1016/j.future.2016.11.009.

[14] Ning Z, Liao J, Zhang F, Shi W. Preliminary study of trusted execution environments on heterogeneous edge platforms. In *Proc. the 2018 IEEE/ACM Symposium on Edge Computing*, Dec. 2018, pp.421-426. DOI: 10.1109/SEC.2018.00057.

[15] Costan V, Devadas S. Intel SGX explained. *IACR Cryptol. ePrint Arch.*, 2016, 2016: Article No. 86.

[16] Kaplan D, Powell J, Woller T. AMD memory encryption. https://developer.amd.com/wordpress/media/2013/12/AMD_Memory_ Encryption_ Whitepaper_v7-Public.pdf, Dec. 2021.

[17] Ngabonziza B, Martin D, Bailey A, Cho H, Martin S. TrustZone explained: Architectural features and use cases. In *Proc. the 2nd IEEE International Conference on Collaboration and Internet Computing*, Nov. 2016, pp.445-451. DOI: 10.1109/CIC.2016.065.

[18] Kim T, Park J, Woo J, Jeon S, Huh J. ShieldStore: Shielded in-memory key-value storage with SGX. In *Proc. the 14th EuroSys Conference 2019*, Mar. 2019, Article No. 14. DOI: 10.1145/3302424.3303951.

[19] Arnautov S, Trach B, Gregor F *et al.* SCONE: Secure Linux containers with intel SGX. In *Proc. the 12th USENIX Symposium on Operating Systems Design and Implementation*, Nov. 2016, pp.689-703.

[20] Priebe C, Vaswani K, Costa M. EnclaveDB: A secure database using SGX. In *Proc. the 2018 IEEE Symposium on Security and Privacy*, May 2018, pp.264-278. DOI: 10.1109/SP.2018.00025.

[21] Tsai C C, Porter D E, Vij M. Graphene-SGX: A practical library OS for unmodified applications on SGX. In *Proc. the 2017 USENIX Annual Technical Conference*, July 2017, pp.645-658.

[22] Barbalace A, Lyerly R, Jelesnianski C, Carno A, Chuang H R, Legout V, Ravindran B. Breaking the boundaries in heterogeneous-ISA datacenters. *ACM SIGARCH Computer Architecture News*, 2017, 45(1): 645-659. DOI: 10.1145/3093337.3037738.

[23] Barbalace A, Sadini M, Ansary S, Jelesnianski C, Ravichandran A, Kendir C, Murray A, Ravindran B. Popcorn: Bridging the programmability gap in heterogeneous-ISA platforms. In *Proc. the 10th European Conference on Computer Systems*, Apr. 2015, Article No. 29. DOI: 10.1145/2741948.2741962.

[24] Gu J, Hua Z, Xia Y, Chen H, Zang B, Guan H, Li J. Secure live migration of SGX enclaves on untrusted cloud. In *Proc. the 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, June 2017, pp.225-236. DOI: 10.1109/DSN.2017.37.

[25] Choy S, Wong B, Simon G, Rosenberg C. The brewing storm in cloud gaming: A measurement study on cloud to end-user latency. In *Proc. the 11th Annual Workshop on Network and Systems Support for Games*, Nov. 2012. DOI: 10.1109/NetGames.2012.6404024.

[26] Furlong M, Quinn A, Flinn J. The case for determinism on the edge. In *Proc. the 2nd USENIX Workshop on Hot Topics in Edge Computing*, July 2019.

[27] Ha K, Abe Y, Eiszler T, Chen Z, Hu W, Amos B, Upadhyaya R, Pillai P, Satyanarayanan M. You can teach elephants to dance: Agile VM handoff for edge computing. In *Proc. the 2nd ACM/IEEE Symposium on Edge Computing*, Oct. 2017, Article No. 12. DOI: 10.1145/3132211.3134453.

[28] Nadgowda S, Suneja S, Bila N, Isci C. Voyager: Complete container state migration. In *Proc. the 37th IEEE International Conference on Distributed Computing Systems*, June 2017, pp.2137-2142. DOI: 10.1109/ICDCS.2017.91.

[29] Jamshidi P, Ahmad A, Pahl C. Cloud migration research: A systematic review. *IEEE Transactions on Cloud Computing*, 2013, 1(2): 142-157. DOI: 10.1109/TCC.2013.10.

[30] Zhu J, Hou R, Wang X *et al.* Enabling rack-scale confidential computing using heterogeneous trusted execution environment. In *Proc. the 2020 IEEE Symposium on Security and Privacy*, May 2020, pp.1450-1465. DOI: 10.1109/SP40000.2020.00054.

[31] Hua Z, Gu J, Xia Y, Chen H, Zang B, Guan H. vTZ: Virtualizing ARM TrustZone. In *Proc. the 26th USENIX Security Symposium*, Aug. 2017, pp.541-556.

[32] Nightingale E B, Hodson O, McIlroy R, Hawblitzel C, Hunt G. Helios: Heterogeneous multiprocessing with satellite kernels. In *Proc. the 22nd ACM SIGOPS Symposium on Operating Systems Principles*, Oct. 2009, pp.221-234. DOI: 10.1145/1629575.1629597.

[33] Piraghaj S F, Dastjerdi A V, Calheiros R N, Buyya R. A framework and algorithm for energy efficient container consolidation in cloud data centers. In *Proc. the 2015 IEEE International Conference on Data Science and Data Intensive Systems*, Dec. 2015, pp.368-375. DOI: 10.1109/DSDIS.2015.67.

[34] Wang H, Shi P, Zhang Y. JointCloud: A cross-cloud cooperation architecture for integrated internet service customization. In *Proc. the 37th IEEE International Conference on Distributed Computing Systems*, June 2017, pp.1846-1855. DOI: 10.1109/ICDCS.2017.237.

[35] Baumann A, Peinado M, Hunt G. Shielding applications from an untrusted cloud with Haven. *ACM Transactions on Computer Systems*, 2015, 33(3): Article No. 8. DOI: 10.1145/2799647.

[36] Hunt T, Zhu Z, Xu Y, Peter S, Witchel E. Ryoan: A distributed sandbox for untrusted computation on secret data. In *Proc. the 12th USENIX Symposium on Operating Systems Design and Implementation*, Nov. 2016, pp.533-549.

[37] Ohrimenko O, Schuster F, Fournet C, Mehta A, Nowozin S, Vaswani K, Costa M. Oblivious multi-party machine learning on trusted processors. In *Proc. the 25th USENIX Conference on Security Symposium*, August 2016, pp.619-636.

[38] Shinde S, Le Tien D, Tople S, Saxena P. Panoply: Low-TCB Linux applications with SGX enclaves. In *Proc. the 24th Annual Network and Distributed System Security Symp.*, Feb. 26-Mar. 1, 2017. DOI: 10.14722/ndss.2017.23500.

[39] Schuster F, Costa M, Fournet C, Gkantsidis C, Peinado M, Mainar-Ruiz G, Russinovich M. VC3: Trustworthy data analytics in the cloud using SGX. In *Proc. the 2015 IEEE Symposium on Security and Privacy*, May 2015, pp.38-54. DOI: 10.1109/SP.2015.10.

[40] Li M, Zhang Y, Lin Z, Solihin Y. Exploiting unprotected I/O operations in AMD's secure encrypted virtualization. In *Proc. the 28th USENIX Security Symposium*, Aug. 2019, pp.1257-1272.

[41] Morbitzer M, Huber M, Horsch J. Extracting secrets from encrypted virtual machines. In *Proc. the 9th ACM Conference on Data and Application Security and Privacy*, Mar. 2019, pp.221-230. DOI: 10.1145/3292006.3300022.

[42] Alves T, Felton D. TrustZone: Integrated hardware and software security. *ARM White Paper*, 2004, 3(4): 18-24.

[43] Sun H, Sun K, Wang Y, Jing J. TrustOTP: Transforming smartphones into secure one-time password tokens. In *Proc. the 22nd ACM SIGSAC Conference on Computer and Communications Security*, Oct. 2015, pp.976-988. DOI: 10.1145/2810103.2813692.

[44] Santos N, Raj H, Saroiu S, Wolman A. Using ARM TrustZone to build a trusted language runtime for mobile applications. In *Proc. the 19th International Conference on Architectural Support for Programming Languages and Operating Systems*, Feb. 2014, pp.67-80. DOI: 10.1145/2541940.2541949.

[45] Zhang N, Sun K, Lou W, Hou Y T. CaSE: Cache-assisted secure execution on ARM processors. In *Proc. the 2016 IEEE Symposium on Security and Privacy*, May 2016, pp.72-90. DOI: 10.1109/SP.2016.13.

[46] Guan L, Liu P, Xing X, Ge X, Zhang S, Yu M, Jaeger T. TrustShadow: Secure execution of unmodified applications with ARM TrustZone. In *Proc. the 15th Annual International Conference on Mobile Systems, Applications, and Services*, June 2017, pp.488-501. DOI: 10.1145/3081333.3081349.

[47] Zhao S, Zhang Q, Qin Y, Feng W, Feng D. SecTEE: A software-based approach to secure enclave architecture using TEE. In *Proc. the 2019 ACM SIGSAC Conference on Computer and Communications Security*, Nov. 2019, pp.1723-1740. DOI: 10.1145/3319535.3363205.

[48] Lind J, Priebe C, Muthukumaran D *et al.* Glamdring: Automatic application partitioning for Intel SGX. In *Proc. the 2017 USENIX Annual Technical Conference*, July 2017, pp.285-298.

[49] Soares L, Stumm M. FlexSC: Flexible system call scheduling with exception-less system calls. In *Proc. the 9th USENIX Conference on Operating Systems Design and Implementation*, Oct. 2010, pp.33-46.

[50] Rott J. Intel® advanced encryption standard instructions (AES-NI). https://www.intel.com/content/www/us/en/developer/articles/technical/advanced-encryption-standard-instructions-aes-ni.html, Dec. 2021.

[51] McCune J M, Li Y, Qu N, Zhou Z, Datta A, Gligor V, Perrig A. TrustVisor: Efficient TCB reduction and attestation. In *Proc. the 2010 IEEE Symposium on Security and Privacy*, May 2010, pp.143-158. DOI: 10.1109/SP.2010.17.

[52] Zhang F, Chen J, Chen H, Zang B. CloudVisor: Retrofitting protection of virtual machines in multi-tenant cloud with nested virtualization. In *Proc. the 23rd ACM Symposium on Operating Systems Principles*, Oct. 2011, pp.203-216. DOI: 10.1145/2043556.2043576.

[53] Dautenhahn N, Kasampalis T, Dietz W, Criswell J, Adve V. Nested kernel: An operating system architecture for intra-kernel privilege separation. *ACM SIGPLAN Notices*, 2015, 50(4): 191-206. DOI: 10.1145/2694344.2694386.

[54] Nelson L, Sigurbjarnarson H, Zhang K, Johnson D, Bornholt J, Torlak E, Wang X. Hyperkernel: Push-button verification of an OS kernel. In *Proc. the 26th Symposium on Operating Systems Principles*, Oct. 2017, pp.252-269. DOI: 10.1145/3132747.3132748.

[55] Klein G, Elphinstone K, Heiser G *et al.* sel4: Formal verification of an OS kernel. In *Proc. the 22nd ACM SIGOPS Symposium on Operating Systems Principles*, Oct. 2009, pp.207-220. DOI: 10.1145/1629575.1629596.

[56] Baumann A. Hardware is the new software. In *Proc. the 16th Workshop on Hot Topics in Operating Systems*, May 2017, pp.132-137. DOI: 10.1145/3102980.3103002.

[57] Ferraiuolo A, Baumann A, Hawblitzel C, Parno B. Komodo: Using verification to disentangle secure-enclave hardware from software. In *Proc. the 26th Symposium on Operating Systems Principles*, Oct. 2017, pp.287-305. DOI: 10.1145/3132747.3132782.

[58] Brasser F, Gens D, Jauernig P, Sadeghi A R, Stapf E. SANCTUARY: ARMing TrustZone with user-space enclaves. In *Proc. the 26th Annual Network and Distributed System Security Symposium*, Feb. 2019. DOI: 10.14722/ndss.2019.23448.

[59] Gu J, Wu X, Zhu B, Xia Y, Zang B, Guan H, Chen H. Enclavisor: A hardware-software co-design for enclaves on untrusted cloud. *IEEE Transactions on Computers*, 2021, 70(10): 1598-1611. DOI: 10.1109/TC.2020.3019704.

[60] Levin R, Cohen E, Corwin W, Pollack F, Wulf W. Policy/mechanism separation in Hydra. In *Proc. the 5th ACM Symposium on Operating Systems Principles*, Nov. 1975, pp.132-140. DOI: 10.1145/800213.806531

[61] Liedtke J. Improving IPC by kernel design. In *Proc. the 14th ACM Symposium on Operating Systems Principles*, Dec. 1993, pp.175-188. DOI: 10.1145/168619.168633.

[62] David F M, Chan E, Carlyle J C, Campbell R H. CuriOS: Improving reliability through operating system structure. In *Proc. the 8th USENIX Conference on Operating Systems Design and Implementation*, Dec. 2008, pp.59-72.

[63] Gu J, Wu X, Li W, Liu N, Mi Z, Xia Y, Chen H. Harmonizing performance and isolation in microkernels with efficient intra-kernel isolation and communication. In *Proc. the 2020 USENIX Annual Technical Conference*, July 2020, pp.401-417.

[64] Hildebrand D. An architectural overview of QNX. In *Proc. the Workshop on Micro-Kernels and Other Kernel Architectures*, Apr. 1992, pp.113-126.

[65] Ji D, Zhang Q, Zhao S, Shi Z, Guan Y. MicroTEE: Designing TEE OS based on the microkernel architecture. In *Proc. the 18th IEEE International Conference on Trust, Security and Privacy in Computing and Communications/13th IEEE International Conference on Big Data Science and Engineering*, Aug. 2019, pp.26-33. DOI: 10.1109/TrustCom/BigDataSE.2019.00014.

[66] DeVuyst M, Venkat A, Tullsen D M. Execution migration in a heterogeneous-ISA chip multiprocessor. In *Proc. the 17th International Conference on Architectural Support for Programming Languages and Operating Systems*, Mar. 2012, pp.261-272. DOI: 10.1145/2150976.2151004.

[67] Gordon M S, Jamshidi D A, Mahlke S, Mao Z M, Chen X. COMET: Code offload by migrating execution transparently. In *Proc. the 10th USENIX Symposium on Operating Systems Design and Implementation*, Oct. 2012, pp.93-106.

Jin-Yu Gu received his B.S. degree in software engineering from Shanghai Jiao Tong University, Shanghai, in 2016. He is now a Ph.D. candidate at the Institute of Parallel and Distributed Systems and the School of Software, Shanghai Jiao Tong University, Shanghai. His research interests include operating systems, computer architecture, and security.

Hao Li received his B.S. degree in software engineering from Shanghai Jiao Tong University, Shanghai, in 2020. He is a Master student at the Institute of Parallel and Distributed Systems and the School of Software, Shanghai Jiao Tong University, Shanghai. His research interests include computer architecture and security.

Yu-Bin Xia received his diploma degree from Software School, Fudan University, Shanghai, in 2004, and his Ph.D. degree in computer science and technology from Peking University, Beijing, in 2010. He is currently an associate professor in Shanghai Jiao Tong University, Shanghai. His research interests include computer architecture, operating system, virtualization, and security.

Hai-Bo Chen received his B.S. and Ph.D. degrees in computer science from Fudan University, Shanghai, in 2004 and 2009, respectively. He is currently a professor and the director of the Institute of Parallel and Distributed Systems, Shanghai Jiao Tong University, Shanghai. He is a distinguished member of both CCF and ACM. His research interests include operating systems, and parallel and distributed systems.

Cheng-Gang Qin received his Ph.D. degree in computer science from Graduate School of the Chinese Academy of Sciences, Beijing, in 2012. He is currently a senior technical expert in the Ant Group, Hangzhou. His research interests include operating system, computer architecture and system security.

Zheng-Yu He received his Ph.D. degree in computer engineering from Georgia Institute of Technology, Atlanda in 2012. He currently leads the technical infrastructure in Ant Group, Hangzhou, and his research interests include operating systems, programming models and software engineering.