

A K M Mubashwir Alam\*, Sagar Sharma, and Keke Chen

# SGX-MR: Regulating Dataflows for Protecting Access Patterns of Data-Intensive SGX Applications

**Abstract:** Intel SGX has been a popular trusted execution environment (TEE) for protecting the integrity and confidentiality of applications running on untrusted platforms such as cloud. However, the access patterns of SGX-based programs can still be observed by adversaries, which may leak important information for successful attacks. Researchers have been experimenting with Oblivious RAM (ORAM) to address the privacy of access patterns. ORAM is a powerful low-level primitive that provides application-agnostic protection for any I/O operations, however, at a high cost. We find that some application-specific access patterns, such as sequential block I/O, do not provide additional information to adversaries. Others, such as sorting, can be replaced with specific oblivious algorithms that are more efficient than ORAM. The challenge is that developers may need to look into all the details of application-specific access patterns to design suitable solutions, which is time-consuming and error-prone. In this paper, we present the lightweight SGX based MapReduce (SGX-MR) approach that regulates the dataflow of data-intensive SGX applications for easier application-level access-pattern analysis and protection. It uses the MapReduce framework to cover a large class of data-intensive applications, and the entire framework can be implemented with a small memory footprint. With this framework, we have examined the stages of data processing, identified the access patterns that need protection, and designed corresponding efficient protection methods. Our experiments show that SGX-MR based applications are much more efficient than the ORAM-based implementations.

**Keywords:** SGX-based data analytics, MapReduce, Data flow regularization, Access Patterns, ORAM

DOI 10.2478/popets-2021-0002

Received 2020-05-31; revised 2020-09-15; accepted 2020-09-16.

---

\*Corresponding Author: A K M Mubashwir Alam:

Marquette University, E-mail: mubashwir.alam@marquette.edu

Sagar Sharma: HP Inc., E-mail: sagar.shamra@hp.com

Keke Chen: Marquette University, E-mail:

keke.chen@marquette.edu

## 1 Introduction

With the development of resource-starving applications in big data, artificial intelligence, and the Internet of Things (IoT), cloud computing has become popular for its storage, computation scalability, and accessibility. However, when uploading data and conducting computations in the cloud, data owners lose full control of their data and must trust that service providers can take care of security well. In practice, due to the wide attack surface of the cloud software stack and the unique co-tenancy feature [25], we have witnessed frequent cloud security incidents. Practitioners using public clouds have experienced at least one cloud-related security incident, and more than 93% of organizations are moderately or extremely concerned about cloud security. The challenge is how to improve cloud security and make data owners more confident in using public clouds.

Researchers have been experimenting with novel crypto approaches, such as fully homomorphic encryption (FHE) [4] and secure multi-party computation (SMC) [13, 20], to address this problem. However, such pure software-based cryptographic solutions are still too expensive to be practical. Their overheads in storage, communication, and computation can be in many orders of magnitudes higher than the plaintext solutions. Recent advances in hybrid protocols, e.g., for machine learning [20, 21, 29], strive to reduce the overall costs of the frameworks by blending multiple crypto primitives to implement different algorithmic components. Although we have seen several close-to-practical solutions [20, 29], they are specifically designed for a particular problem or algorithm and difficult to generalize and optimize for new domains.

During the past five years, the trusted execution environment (TEE) has emerged as a more efficient approach to addressing the performance and usability issues in secure outsourced computation. It provides

hardware support to create an isolated environment within the cloud server that cannot even be compromised by adversaries who control the entire system software stack, including the operating system. Intel Software Guard Extension (SGX) [8] is probably the most well-known TEE hardware and SDK implementation. Since 2015, SGX has been available in most Intel CPUs. Using SGX, a user can run their sensitive computations in a TEE called *enclave*, which uses a hardware-assisted mechanism to preserve the privacy and integrity of enclave memory. With SGX, users can pass encrypted data into the enclave, decrypt it, compute with plain text data, encrypt the result, and return it to the untrusted cloud components.

SGX aims to protect the confidentiality and integrity of the in-enclave data, code, and computation. While it provides much better performance than expensive pure software-based crypto approaches, it leaks interactions and access patterns between the untrusted area and the enclave. Recent studies [10, 27, 36] show that some data access patterns may reveal application-level information to adversaries, even though the data blocks on the untrusted memory are encrypted. Thus, several efforts [1, 27] adopt the concept of oblivious RAM (ORAM) [12] to address the access-pattern leakage problem. While ORAM is a generic solution covering all types of SGX applications by protecting each I/O access, it incurs high costs. With the state of the art ORAM schemes so far [32, 33], each block access still comes with an additional  $O(\log N)$  overhead, with  $N$  numbers of blocks in the pool. Furthermore, for data-intensive applications, the data to be processed is normally larger than the limited physical enclave memory (i.e., the Enclave Page Cache) [8]. Thus, page swapping happens frequently, which can be precisely captured by adversaries via observing page fault interrupts [31] to infer the access pattern of sensitive data.

## 1.1 Scope of Our Research

While ORAM provides a generic primitive for hiding data access patterns for all I/O operations, we find that it is unnecessary to hide sequential, uniform, and completely random I/O patterns, especially for data-intensive applications. For example, sequentially reading or writing a block file does not reveal additional information about the file content. However, the ORAM based I/O interface does not distinguish the different types of application patterns to treat them differently. Developers can certainly carefully examine each step

in their applications to decide to use (or not to use) the ORAM operation. However, it is time-consuming and error-prone. Developers who are not familiar with the access-pattern attacks may accidentally mishandle or simply overlook some I/O accesses. We hypothesize that for many data-intensive algorithms, e.g., data mining algorithms, *regulating the dataflow of the application and focusing on protecting specific access patterns* can be more efficient than using the ORAM primitive for all I/O accesses. We propose a framework: SGX-MR to automatically regulate the data flow of data-intensive applications and protect access patterns without depending on the expensive ORAM primitive or scrupulous examination and design by developers. The proposed approach protects not only the data access patterns from the untrusted memory but also record-level access patterns within the enclave threatened by page-fault attacks [31]. SGX-MR provides all of these benefits while keeping the low-level details transparent to the developers.

Our idea is to take the MapReduce framework [9] to regulate data-intensive SGX application dataflows, and then examine the stages of MapReduce processing to address their access pattern leakages. This approach has several advantages. (1) Since the dataflow in MapReduce processing is fixed for all MapReduce applications, it is more efficient for us to analyze each stage’s access pattern, identify possible access-pattern leakages, and apply more efficient oblivious methods than ORAM. Using this systematic approach, it is possible to cover a large number of potential applications. (2) The MapReduce framework is easy to use and has been applied to numerous data-intensive problems, e.g., data mining [18], machine learning [6], text processing [16], and basic statistics. Researchers and practitioners have accumulated extensive experience [6, 17, 18] for solving most data analytics problems with the MapReduce framework during the past decade. Furthermore, Roy et al. [26] also show that certain types of aggregate reducers with user-defined mapper functions will be sufficient to support a large class of data mining algorithms, including clustering, predictive modeling, and recommenders. SGX-MR strategically uses this finding to facilitate a wide range of data analytics while leveraging SGX’s features in solving the related confidentiality and integrity concerns. Therefore, SGX-MR<sup>1</sup> provides

---

<sup>1</sup> SGX-MR does not support the join operation yet, which will be examined for full performance and privacy guarantee in the future.

a versatile framework/tool for developers to implement wide range of data mining algorithms without having to worry about confidentiality and integrity concerns. Note that it is well known that applications requiring random accesses of data or low latency processing may not match the design purpose of MapReduce. Accordingly, SGX-MR will not be a good candidate for such applications.

We have carefully designed the whole SGX-MR framework<sup>2</sup> to minimize the attack surface [30] and achieve the best performance as possible. Specifically, our research has the following contributions.

- We have implemented the lightweight SGX-MR framework that is flexible to adapt to the restricted enclave memory, data volume, and computation complexity.
- We have carefully studied both the access patterns outside and inside the enclave during each stage of the SGX-MR framework and designed robust mitigation methods to prevent access pattern leakages.
- We have conducted extensive component-wise experiments to understand the cost and performance of the SGX-MR framework with the ORAM based SGX approach. The result shows that SGX-MR can be several times faster than ORAM based solutions.

In the remaining sections, we will first give the background knowledge for our approach (Section 2), then dive in the details of the proposed approach (Section 3), present the evaluation result (Section 5), and described the closely related work (Section 6).

## 2 Preliminary

We will present the related background knowledge before we dive into our approach, including the way SGX manages the enclave memory and related issues, the ORAM approach to addressing the known access-pattern attacks, and a brief introduction to MapReduce framework.

### 2.1 Enclave Memory Management

Due to the hardware implementation cost, Intel SGX only reserves 128 MB of physical memory, known as

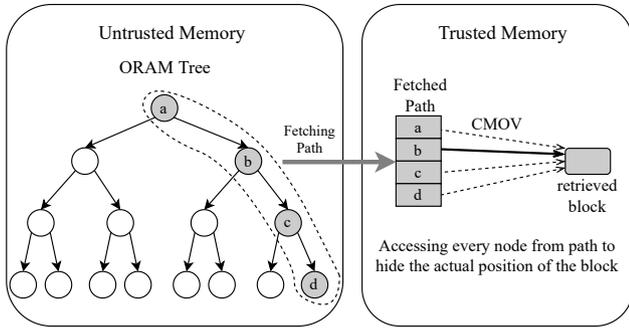
Processor Reserved Memory (PRM). A certain portion of PRM has been used for maintaining SGX’s internal operations. The remaining portion, about 90MB known as the Enclave Page Cache (EPC), can be used by the enclave programs. In the Linux systems, SGX can utilize the Linux virtual memory management for enclave programs, which can thus access a much larger virtual memory address space with contingency for potential page-faults. When a page-fault interrupt happens, the SGX driver will encrypt the selected EPC page and swap it out to the untrusted memory area. This memory management brings two concerns. First, since the SGX driver mostly uses the kernel’s functions for page-fault handling, which an adversarial kernel can manipulate to issue targeted page-fault attack [31]. Shinde et al. [31] have shown such an attack can help adversaries learn the in-enclave access patterns over the EPC pages. Second, It is unknown how efficient the inherent SGX virtual memory management works in applications. Arnavtsov et al. [2] show system-managed EPC page swapping incurs significant costs when handling data larger than the limit of enclave memory. We have also experimentally confirmed this observation. We also designed a block-buffer based application memory management to minimize the memory access pattern. However, due to the large processed data, the application-managed buffer also needs frequent swapping out blocks explicitly, showing no significant performance advantage over implicit system page swapping, as we show in experiments. Thus, for simplicity, a data-intensive application can just use the system virtual memory mechanism, e.g., by setting a sufficiently large enclave heap size. This also gives some flexibility for handling other issues like large objects.

### 2.2 Access-Pattern Based Attacks and ORAM for SGX

When encrypted messages are either accessed from memory or exchanged via the network, an adversary can observe the dataflow and possibly extract sensitive information. Cash et al. [5] and Zhang et al. [35] have shown how an adversary can extract the content of encrypted documents by leveraging access pattern leakages. Ohrimenko et al. [22] have also demonstrated how sensitive information, such as age-group, birthplace, marital status, etc., can be extracted from MapReduce programs by only observing the network flow and memory skew.

---

<sup>2</sup> We will open-source the SGX-MR framework after we finalize the development.



**Fig. 1.** Oblivious block retrieval from the ORAM Tree. The entire path containing the target block is loaded into the enclave to hide the access pattern. Furthermore, the CMOV instruction is used to hide the in-enclave access of the target block to address the page-fault attack.

While SGX can protect the content in the enclave memory from an adversarial OS, it does not provide a systematic mechanism to protect against access pattern leakages. When the program inside the enclave accesses the untrusted memory, the adversarial OS can always monitor the accessed memory addresses. Furthermore, adversaries can modify the virtual memory management mechanism to manipulate the page-fault interrupt generation to figure out in-enclave page access patterns [31].

Consider an example where a pharmaceutical company wants to find top sold items last month by analyzing the daily transactions in the cloud. As sales information is sensitive, they do not want to reveal any statistics from the transactions. Hence in a secure setting, data remains encrypted in untrusted memory and only decrypted in a trusted enclave during computation. A simple way is to use Hashmap to collect the counts of sold items: each item is mapped to a key, which is associated to a key-count pair located at a fixed memory address; each time an item is observed, the corresponding memory address is visited to update the corresponding count. Without access pattern protection, an adversarial OS can observe the memory page access pattern, which can be used to infer the number of a certain item accesses. This access pattern can be used to estimate the number of medicine types, the count of certain medicine, and the total number of sold items.

Researchers have applied the idea of Oblivious RAM (ORAM)[12] to protect the block-level access patterns for SGX applications [1, 27]. ORAM primarily retrieves multiple memory blocks in a path and then randomly shuffles memory blocks periodically to disguise read/write access patterns on encrypted data blocks, where an adversary cannot learn additional information by observing memory access patterns. ORAM was

originally developed for theoretical treatment for software protection[12], and then actively studied for protecting remote access patterns for cloud-based applications [32, 33] to address the privacy concerns with untrusted cloud service providers. By leveraging the ORAM construction, SGX applications can protect the access pattern of untrusted memory from the adversarial OS. The current SGX ORAM methods also partially address the in-enclave page-fault attack on ORAM-related operations. As shown by Figure 1, the popular SGX ORAM solutions, such as ZeroTrace [27] and Obliviate [1], utilize the most efficient Path ORAM [32] or Circuit ORAM [33], which uses a tree structure and paths from the root to the leaves to hide the actual accessed block. The tree and data blocks are maintained in the untrusted area. Once the path containing the target block is identified, it is loaded into the enclave for the enclave program to extract the target block. However, adversaries may monitor page faults to figure out which block in the path is finally accessed. CMOV instruction[23, 24] is a CPU instruction that moves the source operand to destination if a conditional flag is set. Regardless the flag is set or not, the source operand will be read. Therefore, the occurrence of page fault cannot be used to infer whether the source is copied to the destination or not. A simplified example is shown as follows:

```
//if (a < b) x = a else x = b
CMOVL x, a
CMOVGE x, b
```

Consider  $a$  and  $b$  as two sources we want to protect (e.g., the blocks in the ORAM path, as shown in Figure 1). We want to hide the actual access pattern, say, copying  $b$  to the buffer  $x$ . If the comparison result is true, the copy happens with the CMOVL (move if less) line, otherwise with the CMOVGE (move if greater or equal) line. For both lines, the sources  $a$  and  $b$  are read regardless of the condition is true or not. As the ORAM controller’s vital data structures, such as the position map and stash, will be kept inside the enclave, the in-enclave access patterns about the position map and stash may also be under the page-fault attacks. Similar approaches have been used to hide these data-dependent page access patterns. However, as a lower level block I/O interface, SGX ORAM cannot address application-specific page-fault attacks.

## 2.3 MapReduce

MapReduce[9] is a popular programming model and also a processing framework, designed to handle large-scale data with a massively-parallel processing infrastructure. It has several major computation phases: map, optional combiner, shuffle and reduce. The input data is split into fixed-size blocks. Running in parallel, each Mapper takes the user-defined “map” function to process one assigned data block and converts it to key-value pairs. An optional combiner can be used to pre-aggregate the map outputs, if hierarchical aggregation works for the reduce function. Then, all output key-value pairs of mapper (or combiners if they are used) are sorted, grouped, and partitioned individually for Reducers to fetch. Each Reducer then fetches the corresponding share of Map (or Combiner) output in the shuffling phase. After collecting all shares, each Reducer sorts them. Finally, each Reducer applies the user-defined reduce function to process each group of key-value pairs to generate the final result. In both map and reduce phases, multiple processes (Mappers and Reducers) are run in parallel on distributed computing nodes to achieve high throughputs. Noticing the regulated dataflow in MapReduce processing and the extensive experience in developing MapReduce applications accumulated during the past few years [6, 17, 18], we decide to apply this processing framework to the SGX-based data-intensive applications. By addressing the access pattern leakages in this framework, we can protect a large number of applications using our SGX-MR framework.

## 3 The SGX-MR Approach

Starting with an analysis of the typical data flow for data analytics algorithms in the SGX enclave context, we present the SGX-MR framework to integrate MapReduce processing into the SGX enclave environment. Then, we describe the target threat model and analyze possible access pattern leakages under the SGX-MR framework. We will design mitigation methods to address these access pattern leakages.

### 3.1 Features of SGX-Based Data Analytics Algorithms and Threat Model

Typical data analytics applications handle datasets much larger than the enclave memory and often they are sequentially scanned in multiple iterations. Using SGX for data analytics applications has several unique features. For simplicity, datasets are often organized in encrypted blocks and stored in a block-file structure. For security reasons, the SGX enclave program that runs in the protected EPC area cannot access the file system APIs directly. When processed, they are first loaded into the main (untrusted) memory, and then passed to the enclave via SGX APIs. Encrypted data blocks will be decrypted and processed inside the enclave. In this process, there are two security challenges. (1) Adversaries can observe the interactions between the untrusted memory and the enclave. Thus, preserving this access pattern is essential, which has been addressed by ORAM-based approaches, but appears too expensive. (2) Large data processing in enclave will inevitably cause data spill out, either handled implicitly by the system’s page swapping mechanism, or explicitly by applications, e.g., via buffer management. Adversaries can also observe or even manipulate page faults to obtain in-enclave access patterns.

**Threat Model.** Based on the analysis of the features of SGX-based data analytics algorithms, we derive the following threat model. 1) An adversary may compromise the operating system and any hypervisor of the host running the SGX applications. However, it cannot compromise the SGX hardware and enclaves. 2) The adversary may observe all of the data, program executions, and memory management of the system outside the SGX enclaves. It can thus analyze and detect interaction patterns between SGX enclave process and the untrusted memory. 3) A malicious adversary may tamper the application data and program running in the untrusted memory. 4) A powerful adversary may manipulate the data and program execution in the untrusted memory to force page faults within enclave to obtain in-enclave memory access patterns. 5) Denial-of-service and the side-channel attacks based on power analysis, or timing attacks [11], are out of the scope of this paper.

We refer to the basic SGX mechanism [8] for integrity and confidentiality protection. We encrypt the code and data in the untrusted region of our framework with *sgx\_crypto*, a cryptographic library provided by Intel SGX SDK with robust security protection. We design the MapReduce framework and associated algo-

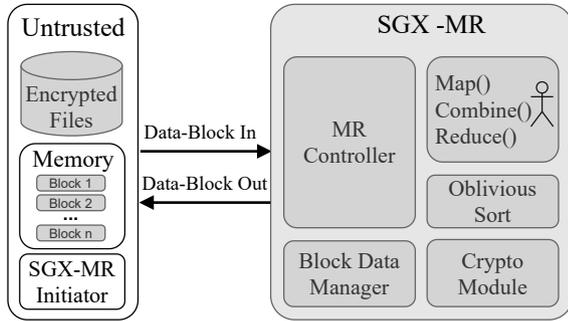


Fig. 2. High-level diagram of SGX-MR: shaded modules or memory either executed in the enclave or remain encrypted.

rhythms in such a manner that the page-fault attacks inside the enclave and the enclave’s access patterns to the untrusted memory are oblivious from the adversary.

### 3.2 Design of SGX-MR

According to the SGX working mechanism and features of data-intensive applications, we partition the entire framework into two parts, i.e., the trusted (enclave) and untrusted parts. Figure 2 shows the components. The untrusted part contains the user-uploaded encrypted data and the small block I/O library. The remaining components of the SGX-MR framework reside in the enclave.

- **Untrusted Part.** Since the data and the I/O library are in the untrusted area, our design needs to address both their confidentiality and integrity. We design a block file structure for encrypted data on disk and in the untrusted memory area. For simplicity, we assume that data records have a fixed size. Both block size and record size are tunable by the user based on the specific application. Each block also contains a message authentication code (MAC) to ensure data integrity. The whole block is encrypted systematically using the AES Counter (CTR) mode functions in the SGX SDK. To minimize the attack surface, we design the library running in the untrusted part to handle only block I/O, and a verification function inside the block data manager in the enclave to capture any adversarial modification on the loaded data blocks.
- **Enclave Part.** The SGX-MR controller handles MapReduce jobs and controls the dataflow of the application. Users only provide the `map()`, `reduce()`, and `combine()` functions to implement the application-specific processing. For simplicity, we will focus on the aggregation-style combine and

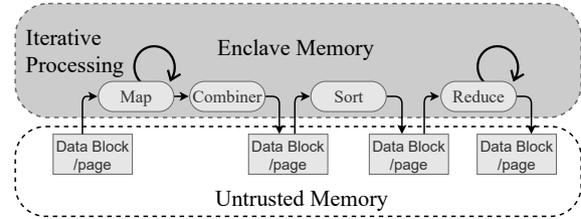


Fig. 3. Regulated dataflows between enclave and main memory

reduce functions, such as COUNT, MAX, MIN, SUM, TOP-K, etc., which have been shown sufficient to handle many data analytics tasks [26]. Remarkably, with our careful design, the binary of the whole SGX-MR framework (without the application-specific map, combine, and reduce functions) takes only about 1.1MB physical memory. With the manually managed memory, we can also work with EPC memory as small as 3–4 blocks. The block size may depend on the specific application (we use block sizes varying from 2 KB to 2 MB for both WordCount and kMeans application in our experiment).

#### 3.2.1 Dataflow Regularization in SGX-MR

With a basic understanding of the components in SGX-MR, we describe how the application dataflow is regulated, which helps simplify access-pattern analysis and protection. While the original MapReduce is designed for parallel processing, Figure 3 sketches how the dataflow regulated by the MapReduce processing pipeline and processed sequentially in SGX-MR, and the interactions between the enclave and the untrusted memory. First of all, input files are processed by the data owner, encoded with the block format via a *file encoding* utility tool, and uploaded to the target machine running the SGX-MR application. Second, within a MapReduce job, all file access requests from the enclave (i.e., mapper reading and reducer writing) have to go through the SGX-MR block I/O module running in the untrusted memory area. Third, the intermediate outputs, e.g., of the Map, Combiner, and Sorting phases, can also be spilled out either by application buffer manager or system’s virtual memory manager, in encrypted form.

Specifically, after the job starts, the Map module will read encrypted data blocks sequentially from the untrusted memory area, decrypt them, and apply the user-provided map function to process the records iter-

atively, which generates the output in key-value pairs. Note that the formats of both the input records and the generated key-value pairs are defined by users for the specific application (for readers who are not familiar with MapReduce programming, please refer to Section 2 and the original paper [9]). The controller accumulates the generated key-value records until they fill up a block, and then sort them by the key. With strictly managed memory, the filled data block will be written to the untrusted area temporarily. If we depend on the virtual memory management, the filled data blocks will stay in the enclave memory and swapped out by the system when needed. We have evaluated both options in experiments and found that application-managed memory has performance advantages. Since we have restricted the reducer functions to a set of hierarchically aggregatable functions (i.e., aggregates can be done locally and then globally), we also design their combiner functions for local aggregation. For example, for the COUNT function, the combiner will generate the local counts for a key, say  $k_i$ :  $((k_i, c_1), (k_i, c_2), \dots, (k_i, c_m))$  if there are  $m$  mappers. The reduce phase will get the final counts of  $\sum_{j=1}^m c_j$ . The inclusion of combiner has two benefits: (1) it can significantly reduce the cost of the most expensive phase: sorting, and (2) it can indirectly address the group size leakage problem in the reduce phase that will be discussed later.

The Combiner outputs will go through the sorting phase, where the sorting algorithm will sort all key-value pairs into groups uniquely identified by the key. The Reduce phase then iteratively processed the groups of key-value pairs identified by the key.

Like the Map phase, key-value pairs stored in blocks will be handled sequentially in the Reduce phase. Specifically, the user-selected *reduce* function from the library takes a group of records with the same key and generates aggregated results. For all the aggregation functions we mentioned earlier, a sequential scan over the group will be sufficient to generate the aggregates. The aggregate of each group is also in the form of key-value pairs, which are accumulated in data blocks, encrypted, and written back to the untrusted area. The above described dataflow keeps the same for all applications that can be cast into the MapReduce processing framework. Now, by specifically addressing the potential access-pattern leakages in the MapReduce data flow, we can effectively protect a broad category of data-intensive SGX applications from access pattern related attacks.

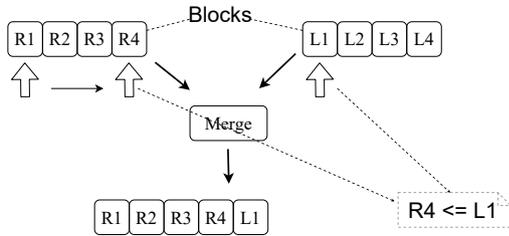
### 3.2.2 Block Design

We fix the record size and records per block to avoid variable length records leaking additional information. Active adversaries can observe the record-level access pattern via page-fault attacks, which allows them to identify record length. Fixed-length records can effectively address this problem, with a cost of padding for short records. One may concern about the efficiency of handling extra-large record that exceeds the limit of the physical enclave memory (e.g., > 100 MB). We admit the padding will take a significant extra cost, but the limit of enclave memory may not be a concern. By adjusting the enclave heap size to sufficiently accommodate blocks, the system’s virtual memory management can handle well — in this case, we prefer not using the application-managed buffer. Experiments have shown the performance difference between different memory management strategies is ignorable.

### 3.2.3 Integrity Guarantee

While SGX assures the integrity of enclave memory, both code and data that reside in untrusted memory remain vulnerable and can be modified. SGX-MR minimizes untrusted execution that was only used for storing and retrieving block data from untrusted memory. The integrity of the untrusted execution will be verified inside the enclave.

We consider three possible attacks to integrity: (1) modify a block, (2) shuffle a block with another block in the same file, and (3) insert a block from a different file (or a phase’s output that is encrypted with the same key). To address all the attacks, we include the following attributes in the block: (i) Block ID, so that block shuffling can be identified, (ii) File Id, so that no block from different files can be inserted, and (iii) the block-level Message Authentication Code (MAC). At the end of each block, a MAC is attached to guarantee the integrity of records, before the whole block is encrypted. We also use the randomized encryption using AES-CTR encoding to make sure the same block will be encrypted to a totally non-distinguishable one so that adversaries cannot trace the generated results in the MapReduce workflow. A simple verification program runs inside the enclave that verifies the IDs and MAC, after reading and decrypting a block.



**Fig. 4.** Access-pattern leakage in MergeSort by observing the interactions between the enclave program and the untrusted memory. If  $R4 \leq L1$ , then the Merge Phase reads R1 to R4 one by one to write sorted output and then reads L1 to L4. By observing this reading pattern (movement of the pointers), the adversary will learn that the entire right sub-list is smaller than the left sub-list.

## 4 Access-Pattern Leakages and Mitigation Methods

Based on the SGX-MR dataflow analysis, we can identify several critical data access patterns pertaining to different phases: Map’s input, intermediate processing, combining, and output; Shuffling/Sorting’s input, sorting, and output; and Reduce’s input, aggregation, and output. Among these access patterns, Map’s input and Reduce’s output involve only sequential block reads/writes. Thus, individually they do not leak additional information except for input and output file sizes (i.e., the number of blocks). However, Reducer’s output, if observed synchronously with its input, may reveal some information as we show in the following. In this section we examine the access pattern leakages in different stages of SGX-MR and discuss the mitigation methods. First, we described leakage in sorting, followed by leakages by page faults, and finally the leakages in the reducing phase of the framework.

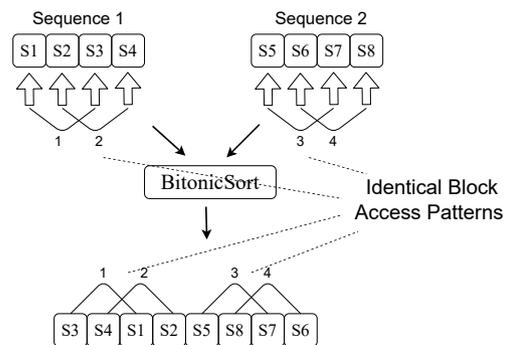
### 4.1 Leakage in Sorting

We start with the most expensive part of the whole data flow - the Sorting phase. Traditionally, the MapReduce framework adopts the MergeSort algorithm for its simplicity. As the shares of Map(or combiner)-output have been sorted individually, MergeSort only needs to merge the sorted shares. However, as we show next, it leaks vital information about the records under sorting.

If we look at the block-wise merge process carefully, we can identify some unusual merging behaviors, which reveals some sensitive attributes of the processed records. Starting from the records in each block sorted

by the Map phase, the main body of MergeSort is to merge and sort two individually sorted lists of records recursively until all records are sorted. Assume the two merged lists contain the records  $\{L_i\}$  and  $\{R_j\}$ , respectively, as shown in Figure 4. If one of the lists has all the values larger (or less) than the other, the corresponding access pattern will be continuously reading the blocks in one list first, followed by the whole list of the other. This pattern may leak sensitive information for real applications. For example, in the WordCount program, adversaries might be able to guess the frequencies of words and derive the word distribution.

**Mitigation Methods.** This access pattern problem involves two parts: block-level pattern in untrusted memory and the page-level pattern in enclave memory. In the following, we address the block-level patterns only with *oblivious sorting* algorithms [3]. The in-enclave page-level patterns will be discussed later. By definition, regardless of the actual values in the sorted list, an oblivious sorting algorithm will take the fixed identical access pattern that is only determined by the size of the list. In the SGX-MR framework, we implement the well-known *BitonicSort* [3] algorithm (See Fig 5. Note that BitonicSort takes  $O(N(\log N)^2)$  block accesses for  $N$  blocks, compared to MergeSort’s cost  $O(N \log N)$ . Our evaluation shows although it is more expensive than MergeSort, the overall cost is still significantly less than MergeSort with ORAM-based oblivious I/O interface.



**Fig. 5.** Uniform block access does not leak any pattern in Bitonic Sort which is oblivious to the data in the blocks. The figure shows the state of the sorting algorithm in the second iteration of sorting

#### 4.1.1 In-Enclave Page-Access-Pattern Leakage

As previously discussed, enclave execution is vulnerable to the page-fault attack [31]. Since the in-enclave pro-

cessing is around records, we look at the related record processing. First, the page-fault attack does not provide additional information for sequential record accesses in executing the map function and the reduce aggregation function. However, we do find some non-sequential in-enclave access patterns that need to be protected from page-fault attacks. They include the in-enclave sorting part of the Sorting phase and the map-output sorting (when a block is filled up) before combining. The repetitive page accesses for the same data block reveal the ordering of records in a pair of blocks, a similar scenario to Figure 4, but happening at the page level.

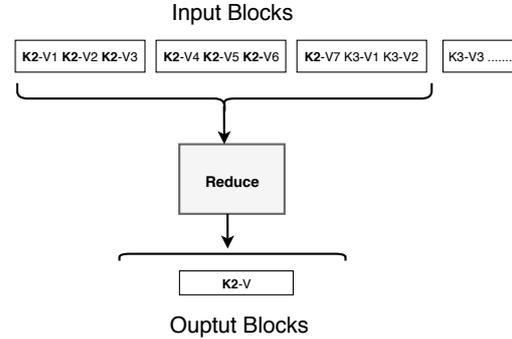
**Mitigation Methods.** At first glance, we can just use the BitonicSort algorithm to hide the record-level access pattern (as a result of page-level access leakage) inside the enclave. However, this is insufficient, as the core operation of this in-enclave BitonicSort, *compare-and-swap*, can still be captured by observing the sequence of page faults and used to reveal the relative ordering between pair of records. The following code snippet shows how possibly the access pattern is associated with the record order.

```
if (a >= b){
// swap a and b, and
// the page access can be observed.
}else{
// no page access.
}
```

To avoid this access-pattern leakage, we adopt the oblivious swap operation applied to different scenarios [7]. The basic idea is to use the CMOV instructions to hide the page access patterns. The oblivious swap operation indeed occurs high cost. We have observed about 2 – 2.5× cost increase for finishing the Sorting phase.

## 4.2 Leakage in Reducing

The Sorting phase will order key-value pairs by the key. As a result, all pairs will be sorted into groups. The controller will sequentially read the sorted blocks and transfer the records containing the same key (i.e., a group of records) to the reduce function. For aggregation-based reduce functions, each group will be reduced to one key-value pair. By observing the input/output ratio, adversaries may estimate group sizes, as shown in Figure 6 (if Combiners are not presented). In extreme cases, multiple blocks can be reduced into one record, which may



**Fig. 6.** Group-wise aggregation in Reduce phase may leak group sizes. K2 spread over three blocks. Thus, the access pattern will be reading three blocks sequentially and then possibly writing out one block.

lead to severe information leakage. For example, it leaks the word frequencies for the WordCount program, which helps estimate the word distribution that, in turn, can be used to guess words.

**Mitigation Methods.** We consider a simple method to address the group-size leakage problem. As described in the previous section, the size difference between the Reducer’s input and output leaks the group sizes. However, as we have made the Combiner component mandatory in SGX-MR thanks to the restricted types of reduce functions, it helps obfuscate the aggregated result in the reduce phase. For example, without Combiners, each record in the reduce phase corresponds to one occurrence of a word (i.e., a key) in the WordCount program, and the group size reveals the frequency of the word. With Combiners, each record entering the Reduce phase may correspond to the pre-aggregate of multiple records. Thus, counting the records in the Reducer input does not give a good estimation of actual group sizes. Note that we do not protect the combining process from page-fault attacks (e.g., the CMOV-based method) for achieving better performance, which will possibly expose partial group sizes. However, adversaries cannot trace back the Reducer input to the Combiner output to utilize the partial group size information because the oblivious sorting phase between them breaks the link entirely. A concern is whether the *frequency* of such combiner-output records still preserves the *ranking* of actual group sizes. The intuition is that the key of a large group may be presented in more map outputs than a key of a small group. We analyzed the observed group size after adding combiners to check whether the size ranking is still preserved. Figure 14 for WordCount and Figure 15 for kMeans both show the ranking of group sizes is not preserved

at all. To completely address the possible ranking leakage, we can further take one of the two candidate approaches: (1) injecting dummy key-value pairs by Combiners, e.g., adding (key, 0) for non-presented keys in the Combiner output for the COUNT function, or (2) disguising the reduce patterns by padding the output block with dummy records, e.g., every time an input block is scanned, output a padded block, regardless of whether or not the actual aggregation happens. In the following, we describe the second method in more detail.

```

load a block;
for each record in the block
  if (the end of the group)
    write the aggregate record
  else
    write a dummy record
write the output block

```

The above process shows the dummy record writing process. Note that all the records will be encrypted and thus non-distinguishable. The concern is that the page-fault attack can catch the execution of the if-else branch, which leaks the information of writing aggregate or dummy record. To address this problem, we also apply the CMOV protection here. Another problem is the increased size of output due to the added dummy records, which can be a significant cost for applications like kMeans that have only a few aggregated records generated. We apply an additional bitonic sorting over the padded blocks and extract top-k elements to get rid of the dummy records, so that the group size ranking is still oblivious to adversaries. In experiments, we will evaluate the additional costs associated with this method.

## 5 Experimental Evaluation

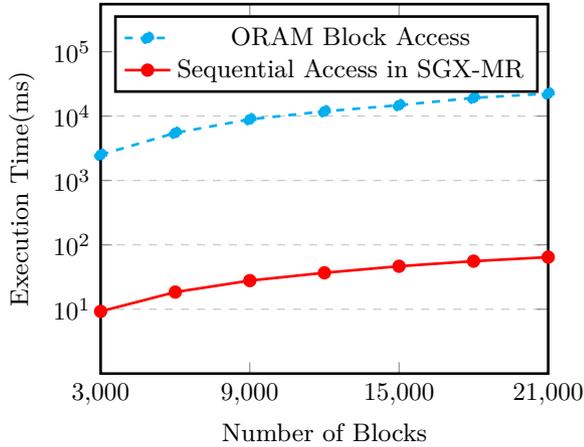
### 5.1 Experiment Setup

SGX-MR is implemented with C++ and the Intel SGX SDK for the Linux environment. Our core framework consists of only about 2000 lines of code. The entire SGX-MR framework runs inside the enclave, except a small component in the untrusted area that handles block-level read/write requests from the enclave. The compiled framework without the application code takes about 1.1MB. To match the design specification, we have implemented a customized bitonic sort that works

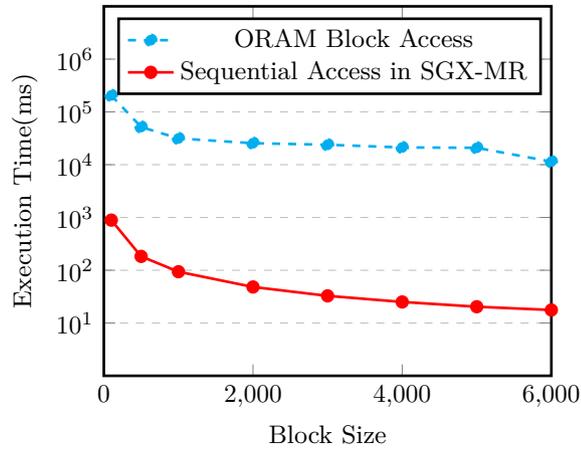
with block-level accesses in the untrusted memory. To protect record level access patterns in the enclave, we also apply the bitonic merge operation to obliviously merge in-enclave blocks. Furthermore, we hide the conditional swaps and sensitive branching pattern by leveraging CMOV instructions. As described in section 4.2, we have implemented mandatory combiners to obfuscate the actual group sizes for the selected aggregators. The experiments were conducted on a Linux machine with an Intel(R) Core(TM) i7-8700K CPU of 3.70GHz processor and 16 GB of DRAM. We use ZeroTrace’s open-source implementation for our experiments. We use 128-bit AES-CTR encryption to encrypt the data blocks in the untrusted memory.

**Sample Applications.** We use two sample applications in our evaluation: WordCount and KMeans. WordCount takes a document collection and output the frequency for each word. It is an essential tool for language modeling and has been included in various tutorials as an example of data-intensive processing. KMeans [14] is a fast, simple clustering algorithm. It takes the initial cluster centroids and iteratively conducts the two steps: (1) cluster membership assignment for each record and (2) the centroid re-computation, until the clustering converges, where the centroids (or all records’ cluster membership) do not change anymore. In all KMeans related experiments, we only execute one iteration of this learning process.

**Baseline Implementation.** To compare SGX-MR with ORAM-based approaches, we have implemented ORAM-based sample applications without depending on the MapReduce framework. In most data analytics programs, aggregating by groups is one of the most expensive core operations. It can be implemented with either hashing or sorting (as in the MapReduce workflow). With ORAM based access-pattern protection, one may wonder whether hashing can be more efficient for enclave-based programs. However, processing hash table in enclave creates page-level pattern [15, 19]. For example, to update a hash value, we must access the position of the particular hash key and update its value list. An adversary can observe this access pattern and possibly figure out the group size that ORAM cannot help. A few recent studies have noticed the problems with enclave based hashtable. Taehoon et al. [15] propose a key-value data-structure for extensive data, but it still leaks the access pattern of the encrypted keys on a hash table. Mishra et al. [19] propose an ORAM based key-value data structure to fully address the access pattern leakage. However, it takes  $O(N)$  complexity to find a key in a worst-case scenario, and it will take  $O(N^2)$  computa-



**Fig. 7.** For sequential access over data ORAM incurs significant cost overhead compared to SGX-MR.

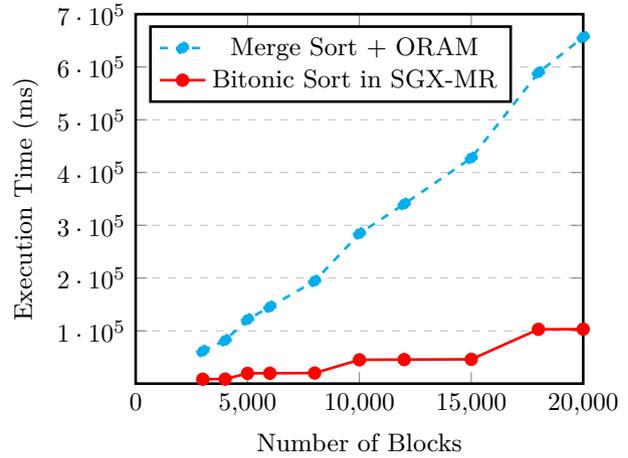


**Fig. 8.** Effect of different block size on ORAM and SGX-MR sequential access, with a fixed file 30 MB.

tional cost to find the aggregated results. On the other hand, MergeSort with ORAM and CMOV takes only  $O(N \log^2 N)$  to sort the records, and a consequent linear scan is sufficient to compute the aggregated results. Thus, in our evaluation, we will take MergeSort with CMOV protection for ORAM based baseline implementations to efficiently address both block and page level attacks.

## 5.2 Performance of Core Operations

**Block Access.** We started our evaluation with the basic operations to understand the sources of performance gain for SGX-MR compared to ORAM. As sequential block accesses do not provide valuable access-pattern information to adversaries, ORAM protection is unnecessary. Figure 7 shows that ORAM protected block ac-

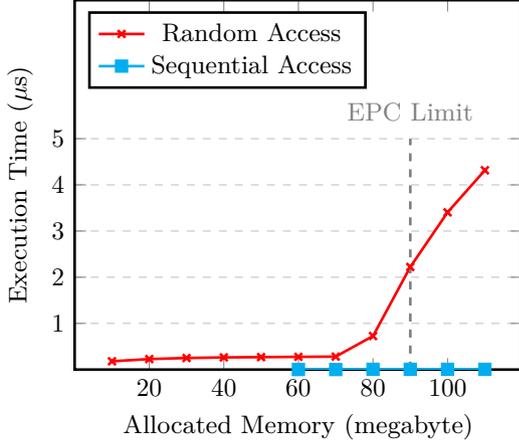


**Fig. 9.** MergeSort with the ORAM block I/O results in significantly higher costs than a dedicated oblivious sorting (e.g., BitonicSort) in SGX-MR. Note: block size = 1 KB.

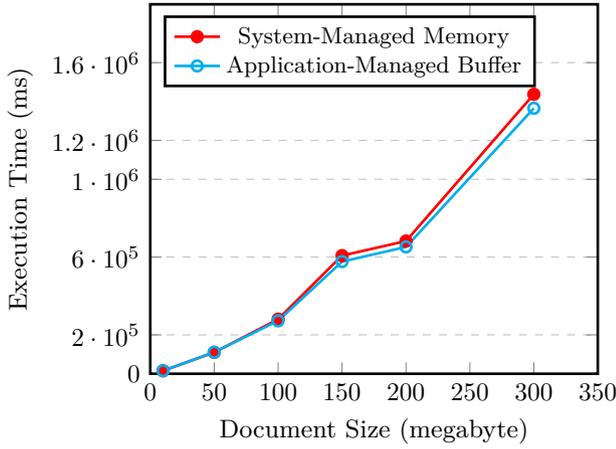
cesses can be 1000x slower than those without protection. As SGX-MR’s input, output, and some intermediate steps, are sequential block accesses, we can utilize this feature to save a significant amount of cost. Since block size is a critical application-specific factor, we measured the impact of block size on performance. Figure 8 shows that a larger block size can significantly reduce the overhead of transferring the same amount of data (30 MB) into enclave. While the block size is increased from 1 KB to 6 MB, the execution time is reduced by about 50 to 100 times for both ORAM and sequential access. Since the ORAM cost is affected by the number of blocks, a larger block size reduces the number of blocks and thus also benefits ORAM.

**Sorting.** Another key operation in our framework is key-value record sorting. We have shown that it is necessary to use oblivious sorting algorithms to protect the important relationship between keys. We use MergeSort in the ORAM setting and BitonicSort for SGX-MR in the comparison. In both solutions, we also used CMOV-based protection for page-fault attacks. Figure 9 shows that BitonicSort under SGX-MR is much faster than MergeSort with ORAM for the block I/O. While the costs of both algorithms are asymptotically the same, i.e.,  $O(n \log^2 n)$ , MergeSort+ORAM appears to have a much high constant factor — it is about five times slower when sorting 20,000 blocks.

**Impact of Enclave Memory Management.** The physical memory allocated for SGX enclaves, called Enclave Page Cache (EPC), is limited to a small size — enclave applications can use only about 90 MB. However, the SGX Linux library can utilize the Linux vir-



**Fig. 10.** When allocated buffer size go beyond EPC limit, system-managed memory incurs significant cost overhead for random data access. Conversely, sequential access has negligible impact on it.



**Fig. 11.** BitonicSort in SGX-MR costs for different memory management strategies. System-managed memory incurs small cost overhead compared to application-managed Buffer.

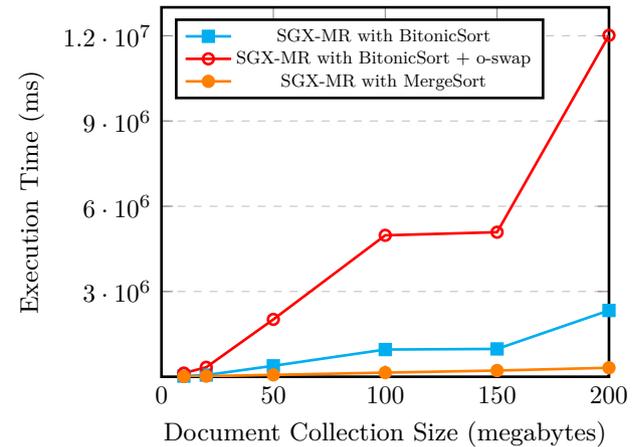
tual memory management so that applications can access the large virtual memory space. Thus, application developers have two options: either entirely depending on the enclave virtual memory management, or manually managing a limited data buffer to minimize the system intervention. Sergei et al. [2] noticed that when the data size is small manually managed buffer can cost less than the enclave virtual memory management as it can reduce the number of unnecessary page swaps. We revisited this issue and show that for sufficiently large datasets (e.g., > 90 MB) these two strategies make not much difference. Figure 10 confirms that page-fault processing takes a significant amount of time. We conduct random block accesses over an array passed to the enclave, which is copied to the enclave heap. When

the array size exceeds the limit of EPC size, the EPC page fault handling mechanism is frequently called for random accesses, which incurs higher costs. However, when working with large data, swapping EPC memory by either buffer management or the system paging does not make much cost different. Figure 11 compares BitonicSort with a four-block buffer management strategy, verses with all blocks loaded into a big array and managed by the system. The cost difference is negligible. However, sometimes the application managed buffer might be desired, e.g., to minimize the memory footprint or better control memory swapping.

### 5.3 Costs of Access-Pattern Protection in SGX-MR

We have discussed several mechanisms to protect the specific access patterns in SGX-MR. In this section, we design a set of experiments to fully understand how these costs contribute to the entire processing costs.

**Oblivious Sorting.** Figure 12 shows the cost comparison with access pattern protection (i.e., BitonicSort) vs. without (i.e., MergeSort) for the WordCount algorithm implemented with SGX-MR. SGX-MR with BitonicSort is approximately seven times slower than SGX-MR with MergeSort. Thus, this cost is significant.



**Fig. 12.** SGX-MR with BitonicSort vs. with MergeSort. O-swap uses CMOV to protect from page-fault attacks. Note: WordCount with block size = 2 MB

#### In-enclave Page-Fault Attack protection.

Next, we evaluate the protection of the page-fault attack on oblivious sorting. Previously, we have shown

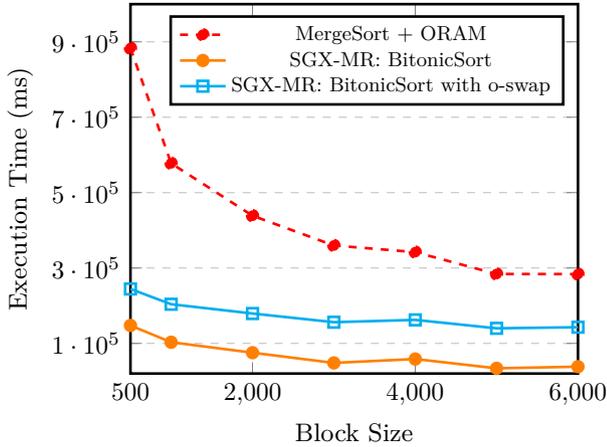


Fig. 13. Effect of block size on the cost of ORAM+MergeSort vs. BitonicSort with/without o-swap.

that protecting access patterns from untrusted memory is not enough. Even in enclave memory, page-level access patterns reveal the order of the records in an oblivious sorting algorithm. As described in section 4, we applied the oblivious swap technique to protect the order of the records within each block.

Fig. 12 and 13 show that, compared to ORAM+MergeSort, the additional cost brought by BitonicSort+o-swap is relatively small.

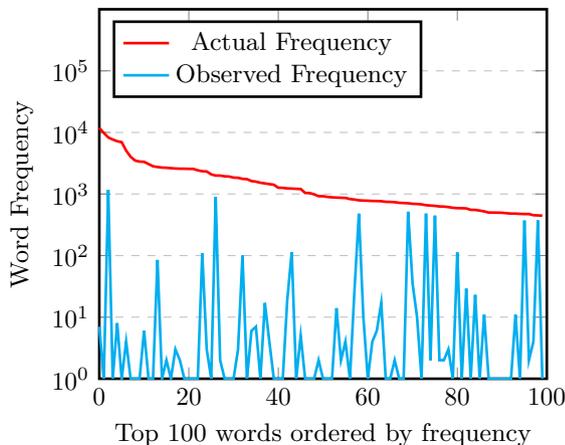


Fig. 14. Observed group sizes before (exactly mapped to word frequencies) and after adding combiners for WordCount.

**Protecting Group Sizes in Reducing.** We have mentioned that by using Combiners mandatorily in our framework we can effectively protect from the group-size-estimation attack in reducing. Figure 14 and 15 how Combiners’ outputs disguise the actual group sizes. It’s well-known that combiners improve performance as they can significantly reduce the number of records

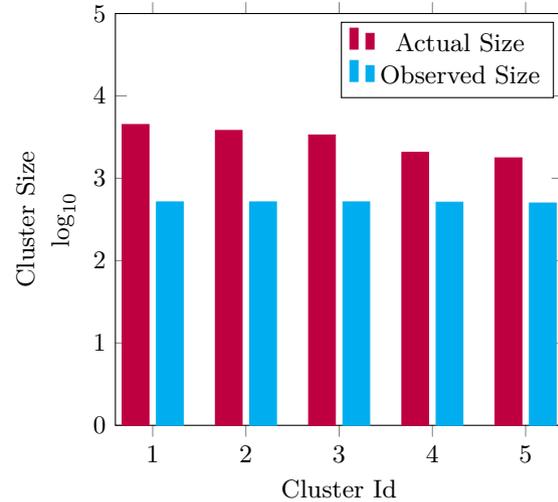


Fig. 15. Observed group sizes before (exactly mapped to the cluster size) and after adding combiners for kMeans.

going to the Reduce phase. We achieved about  $\times 2$  and  $\times 7$  speedup for WordCount and KMeans respectively. As our experiments showed, combiners are especially beneficial for applications with a small number of keys such as KMeans. To further protect the possibly preserved group size ranking, we have also evaluated the dummy-record padding method. Table 1 shows the storage cost of dummy record padding method and its post-processing cost to remove dummy records. The padding will significantly increase the storage cost. However, with the post-processing step, which approximately doubles the processing time of the Reduce phase, the size goes back to the original one, which will not affect the performance of future processing.

## 5.4 Application-Based Evaluation

Finally, we compare the SGX-MR and ORAM implementations for the sample applications. The SGX-MR WordCount compiled code has around 1.6MB and the KMeans code has around 1.5MB. Samples of the Yelp review dataset<sup>3</sup> are used for evaluating the WordCount algorithm. The KMeans algorithm is evaluated with a simulated two-dimensional dataset that fills 30 records per 2 KB block and up to 10,000 blocks. We implement the ORAM-based solutions with the original algorithms (not in the MapReduce style) and use the ZeroTrace block I/O interface for accessing data. As we discussed

<sup>3</sup> <https://www.kaggle.com/yelp-dataset/yelp-dataset>

Table 1. Additional Costs by Padding

Application	Input block	Actual Output block	Padded Output block	w/o post-processing ms	with post-processing ms
KMeans	10,000	1	1,662	33146	51606
WordCount	10,000	3,997	53,439	1022055	2029600

earlier, the ORAM-based baselines use an o-swap protected MergeSort for aggregating the records. Figure 16 and 17 show the overall costs according to the increased data. Overall, the SGX-MR-based implementations for both applications are significantly faster than ORAM-based implementations (Table 2).

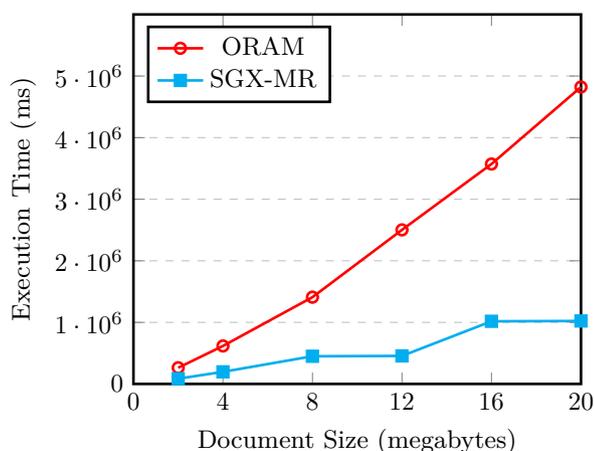


Fig. 16. Application Level Comparison of WordCount Problem

Table 2. Application Level Comparisons of SGX-MR and ORAM for different applications. Block size = 2 KB, data size = 20 MB.

Application	SGX-MR (ms)	ORAM (ms)
KMeans	33078	516634
WordCount	1022974	4821916

## 6 Related Work

There are several pieces of work related to our study. The most popular technique for access-pattern protection would be Oblivious RAM [12]. During the past years, the research on ORAM [32, 33] has been booming due to the rise of cloud computing and the privacy

Application Level Comparison of KMeans Problem

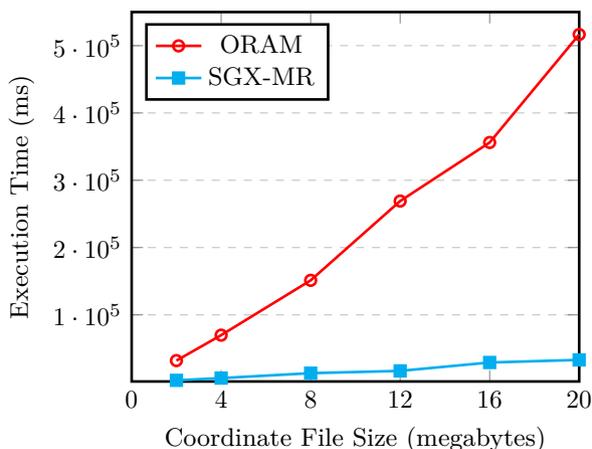


Fig. 17. Application Level Comparison of KMeans Problem

concerns over cloud data store. Tree-based Path ORAM [32] and Circuit ORAM [33] are among the most efficient schemes. ZeroTrace [27] uses both Path and Circuit ORAM to address the access-pattern problem for SGX applications. While using the most efficient implementation of ORAM, we show that ZeroTrace based block I/O still takes a significant overhead over the regular block I/O. Obliviate [1] also uses the Path ORAM scheme to design the I/O interface for an SGX-based secure file system. Similarly, the additional cost is significant. Both schemes try to design a protection mechanism at the block I/O level, providing a high level of transparency to application developers. However, we argue that an application-framework-level protection mechanism can be more efficient.

Researchers also try to extend big data processing platforms to take advantage of SGX. The basic idea is to keep the codebase of current software, such as Hadoop [9] and Spark [34], unchanged as possible, while moving the data-processing parts to the SGX enclave. VC3 [28] applied this strategy for modifying the Hadoop system. It moves the execution of “map” and “reduce” functions to the SGX enclave, while the upper-level functions such as job scheduling and data management still stay outside the enclave. The most part of the Hadoop Java

library is not changed at all. As a result, it achieves the goal of processing encrypted sensitive data in enclaves, but leaves other issues, such as access pattern protection and computation integrity (for the components running in the untrusted memory area), not addressed.

M2R [10] addresses the problem of access-pattern leakage in the shuffling phase of VC3 and proposes to use the oblivious schemes for shuffling. However, other security problems are still not addressed. These top-down approaches have the fundamental problem — unless the whole framework is re-implemented and moved to the enclave, adversaries can easily attack the components running in the untrusted area. Our work is entirely different from these studies. We utilize the MapReduce to regulate the application dataflow so that we can apply access-pattern protection mechanisms to the framework level. Opaque [36] tries to revise Spark for SGX. They also focus on the data access patterns between computing nodes and illustrate how adversaries can utilize these access patterns to infer sensitive information in the encrypted data. To protect the distributed data access patterns, they provide four types of primitive oblivious operators to support the rich Spark processing functionalities: oblivious sort, filter, join, and aggregation. Noticing the problem of computation integrity, they try to move the job controller part, formerly in the master node, to the trusted client-side and design an integrity verification method to detect whether worker nodes process data honestly. However, to reuse the most parts of Spark codebase, it has the most of the system running in the untrusted area, especially for worker nodes. Thus, the local-level integrity guarantee and access-pattern protection might be insufficient. Mostly, it shares the same fundamental problem with VC3 and M2R, as we mentioned.

## 7 Conclusion

ORAM has been a popular solution for protecting access-pattern leakages in SGX-based applications. Yet, it is still very expensive. Furthermore, since ORAM provides a protection mechanism for lower-level block I/O, it cannot address the application-level access-pattern leakage. We also notice that application-specific data access patterns can take advantage of alternative oblivious access techniques to improve performance. Thus, we propose the SGX-MR framework to regulate the dataflow with the MapReduce processing framework, which works for a large class of data-intensive applica-

tions. The regulated dataflow allows us to analyze the access pattern leakages at each stage and develop solutions for all applications implemented with the SGX-MR framework. We have conducted extensive experiments to understand the features of the SGX-MR. The result shows that SGX-MR-based implementations of data-intensive applications work much more efficiently than ORAM-based implementations. We will continue to optimize the phases of the SGX-MR framework to achieve even better performance.

## Acknowledgement

We thank our shepherd Claudio Soriente and all the anonymous reviewers for their guidance and insightful comments on improving this paper. This research is partially supported by NIH (1R43AI136357) and the Wright State Applied Research Corporation.

## References

- [1] A. Ahmad, K. Kim, M. I. Sarfaraz, and B. Lee. Obliviate: A data oblivious file system for Intel SGX. In *the Network and Distributed System Security Symposium*, 2018.
- [2] S. Arnaudov, B. Trach, F. Gregor, T. Knauth, A. Martin, C. Priebe, J. Lind, D. Muthukumar, D. O’Keeffe, M. L. Stillwell, D. Goltzsche, D. Eyers, R. Kapitza, P. Pietzuch, and C. Fetzer. Scone: Secure linux containers with intel sgx. In *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation, OSDI’16*, pages 689–703, Berkeley, CA, USA, 2016. USENIX Association.
- [3] K. E. Batcher. Sorting networks and their applications. In *Proceedings of the April 30–May 2, 1968, Spring Joint Computer Conference, AFIPS ’68 (Spring)*, pages 307–314, New York, NY, USA, 1968. ACM.
- [4] Z. Brakerski and V. Vaikuntanathan. Fully homomorphic encryption from ring-lwe and security for key dependent messages. In *Proceedings of the 31st Annual Conference on Advances in Cryptology, CRYPTO’11*, pages 505–524, Berlin, Heidelberg, 2011. Springer-Verlag.
- [5] D. Cash, P. Grubbs, J. Perry, and T. Ristenpart. Leakage-abuse attacks against searchable encryption. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, CCS ’15*, page 668–679, New York, NY, USA, 2015. Association for Computing Machinery.
- [6] C.-T. Chu, S. K. Kim, Y.-A. Lin, Y. Yu, G. Bradski, and A. Y. Ng. Map-reduce for machine learning on multicore. In *Proceedings Of Neural Information Processing Systems (NIPS)*, 2006.
- [7] S. D. Constable and S. Chapin. libOblivious: A c++ library for oblivious data structures and algorithms. In *Electrical*

- Engineering and Computer Science - Technical Reports*. 184, 2018.
- [8] V. Costan and S. Devadas. Intel sgx explained. *IACR Cryptology ePrint Archive*, 2016:86, 2016.
- [9] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. In *OSDI*, pages 137–150, 2004.
- [10] T. T. A. Dinh, P. Saxena, E. Chang, B. C. Ooi, and C. Zhang. M2R: enabling stronger privacy in mapreduce computation. In *USENIX Security Symposium*, pages 447–462. USENIX Association, 2015.
- [11] H. Gamaarachchi and H. Ganegoda. Power analysis based side channel attack. *CoRR*, abs/1801.00932, 2018.
- [12] O. Goldreich and R. Ostrovsky. Software protection and simulation on oblivious ram. *Journal of the ACM*, 43:431–473, 1996.
- [13] Y. Huang, D. Evans, J. Katz, and L. Malka. Faster secure two-party computation using garbled circuits. In *USENIX Conference on Security*, pages 35–35, 2011.
- [14] A. Jain, M. Murty, and P. Flynn. Data clustering: A review. *ACM Computing Surveys*, 31:264–323, 1999.
- [15] T. Kim, J. Park, J. Woo, S. Jeon, and J. Huh. Shieldstore: Shielded in-memory key-value storage with sgx. In *Proceedings of the Fourteenth EuroSys Conference 2019*, EuroSys '19, New York, NY, USA, 2019. Association for Computing Machinery.
- [16] J. Lin and C. Dyer. *Data-intensive text processing with MapReduce*. Morgan and Claypool Publishers, 2010.
- [17] D. Lyubimov and A. Palumbo. *Apache Mahout: Beyond MapReduce*. CreateSpace Independent Publishing Platform, 2016.
- [18] D. Miner and A. Shook. *MapReduce Design Patterns: Building Effective Algorithms and Analytics for Hadoop and Other Systems*. O'Reilly Media, 2012.
- [19] P. Mishra, R. Poddar, J. Chen, A. Chiesa, and R. A. Popa. Oblix: An efficient oblivious search index. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 279–296, 2018.
- [20] P. Mohassel and Y. Zhang. Secureml: A system for scalable privacy-preserving machine learning. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 19–38, 2017.
- [21] V. Nikolaenko, U. Weinsberg, S. Ioannidis, M. Joye, D. Boneh, and N. Taft. Privacy-preserving ridge regression on hundreds of millions of records. In *IEEE Symposium on Security and Privacy*, pages 334–348, 2013.
- [22] O. Ohrimenko, F. Schuster, C. Fournet, A. Mehta, S. Nowozin, K. Vaswani, and M. Costa. Oblivious multi-party machine learning on trusted processors. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 619–636, Austin, TX, 2016. USENIX Association.
- [23] O. Ohrimenko, F. Schuster, C. Fournet, A. Mehta, S. Nowozin, K. Vaswani, and M. Costa. Oblivious multi-party machine learning on trusted processors. In *25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10-12, 2016.*, pages 619–636, 2016.
- [24] A. Rane, C. Lin, and M. Tiwari. Raccoon: Closing digital side-channels through obfuscated execution. In *Proceedings of the 24th USENIX Conference on Security Symposium, SEC'15*, page 431–446, USA, 2015. USENIX Association.
- [25] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In *Proceedings of the 16th ACM conference on Computer and communications security*, pages 199–212, New York, NY, USA, 2009.
- [26] I. Roy, S. T. V. Setty, A. Kilzer, V. Shmatikov, and E. Witchel. Airavat: Security and privacy for mapreduce. In *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation, NSDI'10*, pages 20–20, Berkeley, CA, USA, 2010. USENIX Association.
- [27] S. Sasy, S. Gorbunov, and C. W. Fletcher. Zerotracer: Oblivious memory primitives from intel SGX. In *25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18-21, 2018*, 2018.
- [28] F. Schuster, M. Costa, C. Fournet, C. Gkantsidis, M. Peinado, G. Mainar-Ruiz, and M. Russinovich. Vc3: Trustworthy data analytics in the cloud using sgx. In *36th IEEE Symposium on Security and Privacy*, 2015.
- [29] S. Sharma and K. Chen. Confidential boosting with random linear classifiers for outsourced user-generated data. In *Computer Security - ESORICS 2019 - 24th European Symposium on Research in Computer Security, Luxembourg, September 23-27, 2019, Proceedings, Part I*, pages 41–65, 2019.
- [30] M.-W. Shih, S. Lee, T. Kim, and M. Peinado. T-sgx: Eradicating controlled-channel attacks against enclave programs. In *Network and Distributed System Security Symposium 2017 (NDSS'17)*. Internet Society, February 2017.
- [31] S. Shinde, Z. L. Chua, V. Narayanan, and P. Saxena. Preventing page faults from telling your secrets. In *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security, ASIACCS16*, page 317–328, New York, NY, USA, 2016. Association for Computing Machinery.
- [32] E. Stefanov, M. V. Dijk, E. Shi, T.-H. H. Chan, C. Fletcher, L. Ren, X. Yu, and S. Devadas. Path oram: An extremely simple oblivious ram protocol. *Journal of the ACM*, 65(4), Apr. 2018.
- [33] X. Wang, H. Chan, and E. Shi. Circuit oram: On tightness of the goldreich-ostrovsky lower bound. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, CCS '15*, page 850–861, New York, NY, USA, 2015. Association for Computing Machinery.
- [34] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica. Spark: Cluster computing with working sets. In *Proceedings of the 2Nd USENIX Conference on Hot Topics in Cloud Computing, HotCloud'10*, pages 10–10, Berkeley, CA, USA, 2010. USENIX Association.
- [35] X. Zhang, G. Li, and J. Feng. Crowdsourced top-k algorithms: An experimental evaluation. *Proc. VLDB Endow.*, 9(8), Apr. 2016.
- [36] W. Zheng, A. Dave, J. G. Beekman, R. A. Popa, J. E. Gonzalez, and I. Stoica. Opaque: An oblivious and encrypted distributed analytics platform. In *USENIX Symposium on Networked Systems Design and Implementation*, 2017.