PRISM – Privacy-Preserving Search in MapReduce

***Abstract*—We present PRISM, a privacy-preserving scheme for word search in cloud computing. Assuming a curious cloud provider, privacy of data stored in the cloud becomes an issue. The main challenge in the context of cloud computing is to design a scheme that achieves privacy while preserving the efficiency of cloud computing. Main approaches like simple encryption, Private Information Retrieval (PIR) and encrypted word search fall short of meeting these requirements. PRISM's idea is to transform the problem of word search into a set of parallel instances of PIR on small datasets. Each PIR instance on a small dataset is efficiently solved by a node in the cloud during the "Map" phase of MapReduce. Outcomes of map computations are then aggregated during the "Reduce" phase. Due to the linearity of PRISM, the simple aggregation of map results yields the final output of the word search operation. We have implemented PRISM on Hadoop MapReduce and evaluated its efficiency using real-world DNS logs. The overhead of PRISM over non-private search is only 11%. Thus, PRISM offers privacy-preserving search that meets cloud computing efficiency requirements. Moreover, PRISM is compatible with standard MapReduce, not requiring any change to the interface or infrastructure.**

I. Introduction

Outsourcing services to clouds has become a major trend in today's IT landscape. Instead of setting up and maintaining their own data centers, cloud users take advantage of *public* clouds, operated by large companies like Google or Amazon. Such public clouds are especially appealing, as they help their customers to lower the cost of ownership while taking advantage of increased scalability, performance, and flexibility. Having started as the most economically viable alternative for small and medium-size enterprises, cloud computing solutions are also being adopted by governmental organizations [12].

Main services offered by clouds are data storage and data processing through user-defined or publicly available programs. Both storage and processing services provided by the cloud are characterized by high performance and quality of service in that most cloud computing services offer reliable storage and management for very large amounts of data as well as highly scalable and parallel data processing facilities. Cloud computing services often rely on specific systems such as Hadoop MapReduce [3], an open source implementation of Google's MapReduce system [20]. Hadoop MapReduce is widely used and public MapReduce clouds are offered to users by companies such as Amazon [1, 28].

The advantages of cloud computing unfortunately come with a high cost in terms of new security and privacy exposures. Apart from classical security challenges of shared services raised by third party intruders or malicious users, such as access control, clustering, service abuse, key-management, and denial-of-service, outsourcing of data storage and processing raises new challenges in the face of potentially malicious cloud providers. *Privacy* of outsourced data appears to be a major requirement in this context.Some regulations are already provisioned as to the privacy protection of outsourced governmental documents [13, 14, 21], e.g., in order to assure privacy against *curious clouds*, clouds with data centers located in "rogue" countries or with insufficient security guarantees, and to avoid data leakage in case of operational failures in the cloud. Along these lines, there is also raising corporate concern about the privacy of sensitive business data stored in the cloud [17]. Although cloud providers thrive to meet the increased privacy demand by certifying their services, cf., Google Apps for Government [36], malicious insiders have still been identified as one of the top threats in cloud computing [18], and users want additional privacy guarantees provided by independent parties.

While encryption of outsourced data by the users seems to be a viable protection mechanism against most privacy problems, classical data encryption mechanisms do not suit the requirements of cloud computing: unlike passive storage devices, the cloud does not only serve as high capacity memory, but is also involved in data processing such as statistical data analysis, log analysis, indexing, data mining, and searching [28]. Although these computations are simple, they have to be carried out in the cloud due to the large amounts of data. However, data processing performed by the cloud would not be feasible with encrypted data. Straightforward download of encrypted data, its decryption and processing in the user environment would also be totally counterproductive by eliminating all performance advantages of cloud computing as highlighted by Chen and Sion [11]. Data has to be processed within the cloud to capitalize on the high performance offered by the cloud.

Among data processing primitives, word search, i.e., verifying, whether a certain word is part of a dataset, is not only one of the most fundamental operations, but surprisingly also one of the most demanded applications in, e.g., MapReduce cloud computing [28]. Related work on search in encrypted data such as Boneh et al. [6], Ogata and Kurosawa [37] falls short of meeting cloud computing privacy and performance requirements. These techniques are designed based on a centralized execution model that is not compatible with today's highly parallel cloud architectures. Even if basic parallelization of related work is feasible, a combination or aggregation of individual results is impossible, therewith shifting computational burden back to the cloud user.

In this paper, we present PRISM, a new scheme for **privacy-preserving and efficient word search** in cloud storage. PRISM pursues two specific objectives: 1.) privacy against potentially malicious cloud providers and 2.) high efficiency through the integration of security mechanisms with the operations performed in the cloud.

In order to achieve efficiency, PRISM takes advantage of the inherent parallelization akin to cloud computing by subdividing the word search problem on a very large encrypted dataset into several instances of word search in small datasets that are executed in parallel. The individual word search operations performed in the cloud yield a result amenable to straightforward *aggregation* in the ultimate phase of the word search operation. The word search operation builds on a Private Information Retrieval (PIR) [32] technique by extending it in order to generate intermediate search results that are still encrypted and that can be *combined* through linear operations to yield the global result of the word search over the entire dataset. Furthermore, in order to achieve high performance and integration with existing cloud systems, PRISM has been designed in order to suit the MapReduce paradigm. During the "Map" phase of MapReduce, individual instances of the word search operation are dispatched to several nodes that compute the intermediate encrypted search results on small datasets. The intermediate results are then combined during the "Reduce" phase into an aggregate thanks to their linearity and the final output of the word search is obtained by the user through the decryption of the aggregate.

PRISM has been implemented on a cloud computing prototype using a standard MapReduce system. PRISM's efficiency has been evaluated through search operations in DNS logs provided by an Internet Service Provider. PRISM's overhead of over search operations without privacy is on average 11% .

The main contributions of PRISM are:

- **Suited to cloud computing:** PRISM is the first privacy-preserving search scheme suited to cloud computing. That is, it brings together storage and search privacy with very high performance by leveraging the efficiency of the MapReduce paradigm. PRISM is not only parallelizable, but also allows efficient combination of individual results.
- **Privacy in the face of potentially malicious cloud providers:** PRISM allows the user to carry out these critical operations in the cloud without even having to trust the cloud provider.
- **Compatible with standard MapReduce:** PRISM only requires a standard MapReduce interface without modifications in the underlying system. PRISM can thus be integrated on any cloud that provides a standard MapReduce interface such as Amazon.
- **Flexible search:** In contrast to classical encrypted keyword search techniques, PRISM is not limited to searching for a fixed set of *predetermined* keywords to be known in advance, but offers flexible search for any words.

It should be noted that even though PIR is an essential building block for PRISM, there is much more to PRISM than a simple parallelization of PIR. PRISM achieves *word search* in cloud storage with a strong emphasis on efficiency, whereas PIR schemes aim at *retrieving specific data* from a central database.

II. Problem Statement and Adversary Model

Throughout this paper, we will use a toy application example to motivate our work. Inspired by recent events [35], we envision a *data retention* scenario. As of today, telecommunication service providers such as Internet service providers or telephone providers, must log and retain (for some time) specific details about clients accessing services. This data retention is required due to regulatory matters in many countries to enable, for example, law enforcement or to combat terrorism. As the sheer amount of data to retain will raise, we expect service providers to outsource their logfiles to clouds. Still, e.g., law enforcement authorities will contact providers and want them to search for words (strings, text, ...) in outsourced files.

Assume service provider $\mathcal{U}$ (the cloud "user") providing DNS services to clients. $\mathcal{U}$ logs each client's access, i.e., $\mathcal{U}$ logs the tuple (timestamp, client ID, hostname queried). Due to the large amount of log data and cost reasons, $\mathcal{U}$ outsources its logfiles into a cloud. Regularly, say each day i, $\mathcal{U}$ creates a new logfile L_i. At the end of a longer period, $\mathcal{U}$ wants (or is forced to) to find out, whether a client was interested in a suspicious host w. So, $\mathcal{U}$ checks, at which day, i.e., in which logfiles L_i, word w occurs. $\mathcal{U}$ queries the cloud for w, and the cloud responses with an answer R telling $\mathcal{U}$ which of the L_i contains w.

Note that $\mathcal{U}$ does not know in advance which word w it has to search for. This automatically disqualifies protocols for *predefined* keyword search in encrypted text, such as PEKS [6] and derivatives. Also note that data retention regulations require outsourced data to be fully recoverable; storing only digests of data in the cloud, e.g., hash values, will not be sufficient.

The cloud is assumed to be untrusted, more precisely semi-honest ("honest-but-curious"). Regulatory matters imply that the cloud must not learn any information about the content it hosts and search queries performed.

Privacy Requirements

We will now formalize privacy requirements for our application. Our application demands for two main types of privacy. *1.) Storage privacy:* The cloud must not be able to infer any details about stored data. *2.) Query privacy:* The cloud must not learn details about $\mathcal{U}$'s queries and results delivered back to $\mathcal{U}$.

1) Storage privacy: The main problem we address with storage privacy is that the cloud (now called "adversary $\mathcal{A}$") must not be able to learn any information about the content. This implies not only the secrecy or confidentiality of the content, but also the inability to compute statistics on the content. For example, if stored logfiles contain the same words, then this must be undetected. While secrecy or confidentiality automatically call for encryption of content before uploading, one (sufficient) approach of additionally preventing adversaries to compute statictics on the content can be the use of semantically secure, ciphertext indistinguishable encryption [26]. More formally, our storage privacy requirements

are captured by IND-CPA cryptographic encryption. As we focus on symmetric encryption (and will discuss why later in Section V-A1), we use the symmetric key IND-CPA definition "IND-P1-C0" introduced by Katz and Yung [31]. Here, public-key encryptions performed by adversary $\mathcal{A}$ are replaced by an encryption oracle that keeps key K of a symmetric encryption E_K secret. We briefly summarize IND-P1-C0: in a *learning phase*, adversary $\mathcal{A}$ can submit a number of plaintexts w_i to an encryption oracle $\mathcal{O}_{E_K}$ that returns ciphertexts $E_K(w_i)$. In the *challenge phase*, $\mathcal{A}$ submits two plaintexts $\{w_0, w_1\}$. Now, $\mathcal{O}_{E_K}$ randomly chooses $b \in \{0, 1\}$ and sends back ciphertext $E_K(w_b)$. Looking at $E_K(w_b)$, $\mathcal{A}$ has to guess b.

Definition 1 (IND-P1-C0 [31])**.** *A cryptographic encryption E_K is IND-P1-C0* **iff** *the probability that $\mathcal{A}$ guesses b correctly is $\leq \frac{1}{2} + \epsilon$, where ϵ is negligible.*

For details, refer to Katz and Yung [31]. Note that the above definition implies secrecy/confidentiality, as we require E_K to be a cryptographic encryption function.

In conclusion, storage privacy requires encryption of content before uploading to the cloud using an IND-P1-C0 encryption mechanism.

2) Query privacy: Interaction between $\mathcal{U}$ and the cloud consists of two steps. First, $\mathcal{U}$ sends $Q(w)$, a search query for word w to the cloud. The cloud processes this query using a protocol $\mathcal{P}$ to produce output $R := \mathcal{P}(Q(x))$. This output is sent back to $\mathcal{U}$. Using this output R and another algorithm $\mathcal{D}$, $\mathcal{U}$ can compute the list of files $\mathcal{D}(R) = \{L_1, L_2, \dots\}$ containing w.

In our scenario, an adversary $\mathcal{A}$ must not learn which word w user $\mathcal{U}$ is looking for. Also, $\mathcal{A}$ must not learn anything about subsequent queries, i.e., search *patterns* should be protected: for example, $\mathcal{A}$ is oblivious of whether the same word is queried for twice. Moreover, $\mathcal{A}$ must not learn which file L_i contains w. That is, $\mathcal{A}$ must not learn for which file the search was "successful" or whether two subsequent queries have produced the same list of files. We capture the above requirements in our definition of query privacy.

Definition 2 (Query Privacy)**.** *A protocol $\mathcal{P}$ to find words on encrypted data provides* query privacy, **iff** *adversary $\mathcal{A}$ with access to ciphertexts $E(w_i)$, a set of queries $Q(w_j)$, a set of all possible words $\mathcal{W}$, $|Q(w_j)| \leq |\mathcal{W}|$, knowledge of protocol $\mathcal{P}$ and outputs R_j (but not $\mathcal{D}$) cannot*

1) *decide whether $\exists Q(w_u), Q(w_v)$ such that $w_u = w_v$.*
2) *decide whether $\exists L_u \in \mathcal{D}(R_u)$, $L_v \in \mathcal{D}(R_v)$ such that $L_u = L_v$.*
3) *discover w_j of any query $Q(w_j)$ or L_u of any result R_u.*

The first two conditions ensure *indistinguishability* of queries and outputs. This avoids linking of any queries or results. The adversary $\mathcal{A}$ cannot decide whether two queries are about the same w, and $\mathcal{A}$ cannot decide whether w was found in the same file on subsequent queries or not. The third condition ensures query and result *confidentiality*. By looking at queries and results, $\mathcal{A}$ can neither discover word w that a query was searching for, nor which file contains w.

We require the total number of queries to be less than the total number of words, $|Q(w_j)| \leq |\mathcal{W}|$, because the adversary could otherwise trivially decide that there is at least two queries for the same word. Note that this requirement is not strictly necessary for the confidentiality property.

In conclusion, the definition of query privacy implies a randomized encryption of queries and a protocol $\mathcal{P}$ that can operate under $\mathcal{A}$'s control on such queries producing "randomized" output.

Limitations: This paper considers only semi-honest clouds. A fully malicious cloud might perform denial-of-service attacks or deviate from proper protocol execution. Similar to "reaction attacks" [29], the cloud might send faulty data back to $\mathcal{U}$, to observe $\mathcal{U}$'s reaction (e.g., sending another query), and to draw some conclusions. Although realistic, we claim that such attacks are out of scope of this paper and leave it for future work. For example, $\mathcal{U}$ might detect faulty responses and take appropriate countermeasures. Moreover, above privacy definitions do not capture many **trivial privacy properties**: for example, the cloud learns the total size of outsourced files or the fact that $\mathcal{U}$ is querying for words. While there exist mitigation strategies, such as padding files, these might be contradictory to cloud efficiency. In addition, our privacy definition does not capture the boundary case where there are more queries than possible words. In such a case, *no system* can provide query privacy. We conjecture that, for many applications, loss of those "trivial" privacy properties can be acceptable. Also, this paper only focuses on curious clouds as the adversary and does not consider, e.g., malicious users. To cope with malicious users, any real-world cloud would require additional security mechanisms such as access control, segregation of users etc [18]. This is out of focus of the paper at hand.

III. Background

A. MapReduce

We target a system suited for the MapReduce paradigm. We will now give a condensed overview of MapReduce, only focusing on aspects necessary to understand PRISM. More details can be found in [3].

Upload. Roughly speaking, a MapReduce cloud comprises a set of "slave" node computers and a "master" computer. While user $\mathcal{U}$ uploads files into the MapReduce cloud, each file is automatically split into blocks, so called *InputSplits*. InputSplits have a fixed size $S_{\text{InputSplit}}$ which is a pre-configured system parameter. If S_{File} denotes the size of an uploaded file, the number of InputSplits c computes to $c = \frac{S_{\text{File}}}{S_{\text{InputSplit}}}$. For each of the c InputSplits, a workload sharing algorithm running on the master selects a slave node and places the InputSplit on it. Also, for fault tolerance, each InputSplit is replicated a number of times and placed on additional slaves.

In addition to data, the MapReduce framework also allows $\mathcal{U}$ to upload "operations", i.e., compiled Java classes. These

classes essentially represent the implementation of three functions.

1) $Scan(\text{INPUTSPLIT}) \rightarrow [(k, v)]$), a functions that takes an InputSplit as an input, parses it, i.e., scans it and generates a set of key-value pairs $[(k, v)]$ out of it.
2) $Map(k, v) \rightarrow [(k', v')]$, a function that takes as an input a single key-value pair (k, v) and outputs a set of "intermediate" key-value pairs $[(k', v')]$.
3) $Reduce([(k', v')]) \rightarrow \text{FILE}$, a function that takes as an input a set of intermediate key-value pairs $[(k', v')]$ and writes arbitrary output into a file.

Uploaded Java classes are replicated and sent to all slave nodes storing an InputSplit.

Map Phase. After data and implementations have been uploaded, $\mathcal{U}$ can specify a single uploaded file (now split into InputSplits) and trigger MapReduce operations on that file. As the name suggests, the first phase of operation is the "Map" phase. Each slave node becomes a "mapper" node. First, each mapper executes $\mathcal{U}$'s $Scan$ function on the InputSplit it stores locally. This generates a set of key-value pairs on each mapper. Furthermore, the mapper node executes $\mathcal{U}$'s Map function on this generated key-value pairs to produce a set of intermediate key-value pairs.

Note that if the cloud consists of at least c slave nodes (more precisely, at least c CPUs), then all InputSplits can be processed in parallel. Otherwise, some slave nodes have to work on multiple InputSplits consecutively, becoming a mapper more than once.

Reduce Phase. Eventually, MapReduce starts the "Reduce" phase. Slave nodes in the MapReduce cloud are scheduled to become "reducers". For each of the generated intermediate pairs (k', v'), MapReduce selects a reducer and sends (k', v') to this reducer. More precisely, MapReduce selects the *same* reducer for all intermediate key-value pairs (k', v') having the same key. So, each reducer receives a set of intermediate key-value pairs $(k'_1, v'_1), (k'_2, v'_2), \ldots$ with identical keys $k'_1 = k'_2 = \ldots$. Each reducer executes $\mathcal{U}$'s reduce function on its set of intermediate key-value pairs and writes the output to a file. This file is sent to $\mathcal{U}$.

After completion of such a MapReduce job on a single uploaded file, $\mathcal{U}$ could invoke MapReduce on another uploaded file. However, for increased efficiency, instead of sequentially processing files, $\mathcal{U}$ can also specify a *set* of uploaded files. MapReduce will process the set of files completely in parallel. Still, Scan/Map/Reduce operations on InputSplits belonging to different files are separated to keep the same semantics as files were processed sequentially.

B. Trapdoor Group Private Information Retrieval

PIR allows a user to retrieve data from a server without revealing which data is retrieved. For PRISM, we make use of a simple and efficient PIR mechanism as previously suggested by Trostle and Parrish [44]. As this PIR mechanism is just a (exchangeable) building block for PRISM, we will only give a conclusive summary of its mode of operation and rationale.

Overview: Matrix $\mathcal{M}$ is a $t \times t$ matrix of elements in $\mathbb{Z}_N$ stored at a server. For example, $N = 2$ for a binary matrix. User $\mathcal{U}$ is specifically interested in the elements of the k^{th} row in $\mathcal{M}$, but the server must not determine $\mathcal{U}$'s interest.

The idea is now that $\mathcal{U}$ sends two "types" of values to the server. For each row that $\mathcal{U}$ is not interested in, he sends a value of the "first" type. For the one row that $\mathcal{U}$ is interested in, he sends a single value of the "second" type. To prevent the server from distinguishing between the two types of values, $\mathcal{U}$ blinds each value with a blinding factor b. This blinding factor can later be removed by $\mathcal{U}$. The server now performs simple additions with received values and elements stored in $\mathcal{M}$. The result is sent back to $\mathcal{U}$ who removes the blinding factor and determines the values of the row of his interest.

Preparation: Assume $\mathcal{U}$ is interested in row k. $\mathcal{U}$ chooses a group $\mathbb{Z}_p$, where p is a prime of m bits. $\mathcal{U}$ also chooses a random $b \in \mathbb{Z}_p$ and t random values $a_i \in \mathbb{Z}_p$. Therewith, $\mathcal{U}$ computes t values $e_i < \frac{p}{t \cdot (N-1)}$ such that: $e_k := 1 + a_k \cdot N$ **and** $\forall i \neq k : e_i := a_i \cdot N$.

Finally, $\mathcal{U}$ computes $\alpha_i := b \cdot e_i \mod p$ and sends the α_i to the server. Other values $(p, m, b, \{e_i\}, \{a_i\})$ remain secret. The server treats α_i as large integers and performs the following integer operations, i.e., without any modulo.

Server computation: Let $\vec{u}$ be the vector $\vec{u} := (\alpha_1, \ldots, \alpha_t)$. The server computes the matrix product $\vec{v}$ and sends it back to $\mathcal{U}$,

$$\begin{aligned} \vec{v} := (\beta_1, \ldots, \beta_t) &= \vec{u} \cdot \mathcal{M} \\ &= (\sum_{i:=1}^{t} \alpha_i \cdot \mathcal{M}_{i,1}, \ldots, \sum_{i:=1}^{t} \alpha_i \cdot \mathcal{M}_{i,t}). \end{aligned}$$

Result analysis: Upon receipt, in order to "un-blind" values, $\mathcal{U}$ computes the t inverse values $z_i := \beta_i \cdot b^{-1} \mod p$. Now, U can conclude that $z_i \mod N$ equals the i^{th} element of the k^{th} row in $\mathcal{M}$. Therewith, $\mathcal{U}$ has retrieved the t elements of the k^{th} row of $\mathcal{M}$ in a privacy-preserving fashion. Note the linearity for two β_i and β'_i received during different PIR runs: $\beta_i \cdot b^{-1} + \beta'_i \cdot b^{-1} = (\beta_i + \beta'_i) \cdot b^{-1} \mod p$. That is, the sum of two individually un-blinded vectors equals un-blinding the sum of two received vectors $\vec{v}, \vec{v'}$. We will later use this linearity during PRISM's reduce phase.

Security rationale: Security and privacy of this protocol are based on the trapdoor group assumption. With only knowledge of α_i, but not secret trapdoor p, it is computationally hard for the server to infer any information about low order bits, i.e., the modulo of z or e_i, cf., Trostle and Parrish [44].

Discussion: Again, we stress that this particular PIR scheme is a building block. It can be exchanged by any PIR scheme where multiple server responses $\vec{v}$ can be combined with a group operation. For example, *any* of the "traditional" PIR protocols based on group homomorphic encryption [32] can be easily adapted for use within PRISM, see Ostrovsky and Skeith [39] for an overview. We have chosen Trostle and Parrish [44] only due to the straightforward way to implement it in Java (Section VI). Using different PIR schemes could

lead to a reduction in the (already small) overhead, but this is out of scope of this paper.

IV. PRISM Protocol

Before going into PRISM's technical details, we will first highlight its main design rationales.

A. Overview

PRISM comprises three parts: *upload* of data into the cloud, the actual MapReduce *search*, and finally the *result analysis* where the user will decide whether the word has been found.

1.) Upload. During upload, user $\mathcal{U}$, encrypts words using symmetric encryption. Ciphertexts are stored in a file, and this file is sent to the MapReduce cloud. In PRISM, we use standard encryption, e.g., a blockcipher like AES, to perform ciphering of data. However, to ensure IND-P1-C0, plaintext is modified before encryption. Therewith, $\mathcal{U}$ can still search for some word w, but the cloud cannot compute statistics about ciphertexts.

2.) Search. Eventually, $\mathcal{U}$ wants to search his encrypted files for some word w. Therefore, $\mathcal{U}$ sends implementations of "algorithms" for the map and reduce phases to the MapReduce cloud, and the cloud executes these on uploaded data. For example, $\mathcal{U}$ sends Java ".class" files for the mappers and Java ".class" files for the reducers. MapReduce distributes these implementations to each mapper and reducer, respectively. PRISM's rationale is to *transform* the word search problem into a set of small PIR instances. To do so, each mapper, scanning through its locally stored InputSplit, creates a binary matrix. Ciphertexts in the InputSplit are assigned to individual elements in that matrix. If a ciphertext is present in an InputSplit, its corresponding element in the matrix is set to either "0" or "1". Using private information retrieval techniques, PRISM can extract the value of a single element in the matrix with the mapper being totally oblivious to which element is extracted. Consequently, $\mathcal{U}$ can specify which element to extract in a privacy-preserving way. All mappers send their obliviously extracted elements as key-value pairs to reducers. Reducers simply sum up received values and return sums to $\mathcal{U}$. Therewith, neither mappers nor receivers can learn any information about which ciphertext $\mathcal{U}$ was interested in.

3.) Result analysis. Finally, $\mathcal{U}$ receives an encrypted sum for each of the originally uploaded files from reducers. $\mathcal{U}$ can decrypt them and decide which of the files contain w. However, due to the probability of "collisions" in matrices, i.e., two different ciphertexts can be assigned to the same element, and due to ambiguities of received sums, $\mathcal{U}$'s decision whether w is inside some file might be wrong. Therefore, PRISM repeats the above process in a total of q so called "rounds". In each of the rounds, a new matrix is generated, elements are set to "1" or "0" depending on the round number, and results are returned as described. This reduces the probability of $\mathcal{U}$ making incorrect decisions.

B. Upload Details

In our toy scenario, a webserver continuously logs access to the cloud. Each day the webserver starts using a new logfile. For simplicity, we assume that entries logged by webservers are simple words.

Overview. One privacy requirement as stated earlier in Section II is that the cloud must not be able derive any details about stored plaintext. We will cover this by using a *stateful cipher* and show its IND-P1-C0 property later in Section V-A.

Definition 3 (Stateful Cipher)**.** *Given a standard symmetric encryption cipher E with secret key K, e.g., AES, we extend E to a* stateful cipher*, by adding "counters" γ_{w_i} that count the history of different inputs w_i. Each time E encrypts w_i, counter γ_{w_i} is increased by one.*

In conclusion, a stateful cipher is a cipher that knows how often it has encrypted a specific plaintext. We will now present a trivial stateful cipher construction and use this in PRISM to encrypt before uploading.

Stateful Cipher Example (see Algorithm 1): User $\mathcal{U}$ and webserver W share a secret key K. For each day d, webserver W computes his key-of-the-day $K_d := \mathrm{HMAC}_K(d)$. W executes Algorithm 1.

For each day, W maintains a hash table containing the list of counters γ_{w_i} in W's local storage. At the beginning of each day, W initializes all counters to 0, i.e., $\gamma_{w_i} = 0$. Now, for each logentry w_i that should be stored in the cloud, W computes γ_{w_i} and increases γ_{w_i}. Then, W computes ciphertext C_i. Webserver W sends ciphertext C_i to the cloud that stores it in this day's file.

Input : words w_i
Output: ciphertexts C_i uploaded to cloud
1 Initialize all γ to 0;
2 **foreach** *word* w_i **do**
3 $\quad \gamma_{w_i} :=$ get (w_i); `//from hash table`
4 $\quad \gamma_{w_i} := \gamma_{w_i} + 1$;
5 $\quad$ insert (w_i, γ_{w_i}); `//into hash table`
6 $\quad C_i := E_{K_d}(w_i, \gamma_{w_i})$;
7 $\quad$ **upload** C_i ;
8 **end**

Algorithm 1: "Stateful Cipher" encryption and upload to MapReduce. Key-of-the-day is K_d.

For the encryption $E_{K_d}(w_i, \gamma_{w_i})$, "," denotes an unambiguous pairing of inputs. We discuss the reason for using a "stateful cipher" over using, e.g., a CBC mode of encryption in Section V-A.

C. Map Phase Details

User $\mathcal{U}$ wants to search a set of files for word w within a period (a set of days) of time. For ease of understanding, we will restrict our description below to PRISM working on a single file specified by the user, i.e., the file of day d.

$\mathcal{U}$ sends map and reduce implementations of PRISM to MapReduce, and the map phase starts. In the following, we describe the PRISM algorithms for, first, the mappers and in Section IV-D the reducers. We would like to stress that the PRISM algorithms, e.g., Java ".class" files, are not encrypted

and not specially protected against a curious cloud. Even though mappers and reducers know what operations they perform, they cannot deduce any private information about stored data or details about the search.

Overview. Before scanning through its local InputSplit, a mapper node creates a matrix where all elements are initialized to "0". PRISM's main idea is that while the mapper scans the ciphertexts in its InputSplit, each ciphertext is assigned to one position, a certain element in the matrix by computing a hash of the ciphertext. Additionally, for each ciphertext, the mapper computes a single bit hash, and if the hash output bit is "1", the mapper puts a "1" in the matrix at the assigned position. The idea is that user $\mathcal{U}$ can also compute the position in the matrix and the one bit hash output for a word w he is looking for. Roughly speaking, $\mathcal{U}$ now queries the mapper for the value of that bit in the matrix using private information retrieval. If the bit retrieved from the mapper differs from the bit computed by $\mathcal{U}$, then $\mathcal{U}$ can decide, e.g., that w is not in this InputSplit.

Problem is that due to the limited size of the matrix and the properties of the hash function, there might be collisions in the assignment process. That is, by chance there can be two different ciphertexts being assigned with the same position in the matrix. By chance, the bit retrieved by $\mathcal{U}$ can therefore be unrelated to w. This problem is amplified by the fact that $\mathcal{U}$ does not only receive a single bit for a single InputSplit, but a combination (the *sum*) of all bits from all mappers working on InputSplits. To mitigate this problem, PRISM repeats generation and filling of matrices a total of q rounds. Also, setting an element in a matrix to "1" depends on the round number. After q rounds, the probability that the information $\mathcal{U}$ retrieved from this mapper is unrelated to w therefore decreases, and $\mathcal{U}$ can finally decide whether w is inside this file.

Definition 4 (PIR Matrix)**.** *A binary matrix $\mathcal{M}$ of size $t \times t$ is called a PIR matrix. This is the matrix used by the mapper to implicitly perform the word search in a privacy-preserving way.*

Definition 5 (Candidate Position)**.** *For each ciphertext C_i in an InputSplit, the candidate position $(\mathcal{X}_i, \mathcal{Y}_i)$ of C_i in $\mathcal{M}$ is computed by $(\mathcal{X}_i||\mathcal{Y}_i) := \lfloor C_i \rfloor_{2\cdot\log_2(t)}$. Here, $\lfloor \dots \rfloor_{2\cdot\log_2(t)}$ denotes truncation after $2\cdot\log_2(t)$ bits. So, the first $\log_2 t$ bits of C_i determine $\mathcal{X}_i$, and the second $\log_2 t$ bits determine $\mathcal{Y}_i$.*

Definition 6 (PIR Input)**.** *If $\mathcal{U}$ is interested in a specific element $(\mathcal{X}, \mathcal{Y})$ in $\mathcal{M}$, he computes PIR input $\{\alpha_1, \alpha_2, \dots, \alpha_t\}$, where $\alpha_{\mathcal{X}} := b \cdot (1 + a_{\mathcal{X}} \cdot N) \mod p$, and $\forall i \neq \mathcal{X}, \alpha_i := b \cdot (a_i \cdot N) \mod p$. Random values b and a_i as well as parameters N and p are chosen as for the Trapdoor Group PIR scheme presented in Section III-B.*

Definition 7 (Column Sum)**.** *The column sum σ_i of the i^{th} column of PIR matrix $\mathcal{M}$ is defined as*

$$\sigma_i := \sum_{\mathcal{M}_{1\leq j\leq t,i}=1} \alpha_j,$$

where $\mathcal{M}_{1\leq j\leq t,i} = 1$ denotes the entries in the i^{th} column of $\mathcal{M}$ that are set to 1.

Note that additions in this definition are integer additions.

The above computation of column sums is simply a digest of the PIR technique by Trostle and Parrish [44]. In short, if a mapper computes such a column sum on a given PIR matrix $\mathcal{M}$ and given PIR inputs α_i, it is impossible for the mapper to derive $(\mathcal{X}, \mathcal{Y})$. $\mathcal{U}$, however, can compute whether $\mathcal{M}_{\mathcal{X},\mathcal{Y}} = 1$, because $\mathcal{M}_{\mathcal{X},\mathcal{Y}} = 1$ **iff** $(\sigma_{\mathcal{Y}} \cdot b^{-1} \mod p) \mod 2 = 1$ holds.

It is important to point out that not only a mapper can compute a candidate position for some ciphertext in its InputSplit, but also $\mathcal{U}$ can compute candidate positions. More precisely, as $\mathcal{U}$ is looking for w, he can compute $E(w, 1)$ and candidate position $(\mathcal{X}_i||\mathcal{Y}_i) := \lfloor E(w, 1) \rfloor_{2\cdot\log_2(t)}$. If w has been uploaded into a particular InputSplit at least once, then this InputSplit contains at least $E(w, 1)$ (maybe also $E(w, 2)$, $E(w, 3), \dots$). Therefore, it is sufficient for $\mathcal{U}$ to search for $E(w, 1)$. We will now give detailed descriptions of $\mathcal{U}$'s Map and Reduce algorithms.

1) PRISM Map Description: To start, $\mathcal{U}$ chooses two system parameters $t, q \in \mathbb{N}$ for PRISM. Parameter t is a security parameter, and q denotes the number of rounds. For day d that $\mathcal{U}$ wants to search for w, he determines the key-of-the-day $K_d := \mathrm{HMAC}_K(d)$ and the target candidate position $(\mathcal{X}_d||\mathcal{Y}_d) := \lfloor E_{K_d}(w, 1) \rfloor_{2\cdot\log_2(2)}$. To prepare PIR, $\mathcal{U}$ computes for day d the t PIR Inputs $\{\alpha_1, \alpha_2, \dots, \alpha_t\}$ as described above. $\mathcal{U}$ sends all α as part of the following map algorithm implementation to the cloud.

All mappers process PRISM in parallel, each of them on its own, locally stored InputSplit of the file corresponding to day d. More precisely, a mapper executes Algorithm 2. Initially, the mapper generates q PIR matrices M_l, where each element is initially set to 0. We will now write $\mathcal{M}_{l,\mathcal{X},\mathcal{Y}}$ to denote element $(\mathcal{X}, \mathcal{Y})$ in matrix $\mathcal{M}_l$.

The mapper node *scans* its local InputSplit consisting of ciphertexts $\{C_1, \dots, C_n\}$. For each ciphertext C_i, the mapper creates a key-value pair (i, C_i). Then, the mapper fills matrices $\mathcal{M}_l$. For pair (i, C_i),

- the mapper computes candidate position $(\mathcal{X}_i||\mathcal{Y}_i) := \lfloor C_i \rfloor_{2\cdot\log_2(t)}$.
- the mapper puts in PIR matrix $\mathcal{M}_j$, in element $\mathcal{M}_{j,\mathcal{X}_i,\mathcal{Y}_i}$, a "1", if the bit $bit_j := \lfloor h(C_i, j) \rfloor_1 = 1$. Here, h denotes a cryptographic hash function and "," again an unambiguous pairing of inputs. If $bit_j = 0$, element $\mathcal{M}_{j,\mathcal{X}_i,\mathcal{Y}_i}$ remains untouched. This means that entries in $\mathcal{M}_j$ can flip from 0 to 1, but never from 1 back to 0.

$\mathcal{X}\backslash\mathcal{Y}$	1	2	$\cdots$	t
1	1	0	$\cdots$	1
2	0	**1**	$\cdots$	0
$\vdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$
t	1	0	$\cdots$	0

Fig. 1. Sample matrix $\mathcal{M}_j$.

Fig. 1 shows a sample $t \times t$ PIR matrix $\mathcal{M}_j$ representing

round j. Initially, all elements have been set to 0. While scanning through the local InputSplit, i.e., the ciphertexts $C_1, \ldots, C_n$, the mapper sets some elements to 1. For example, ciphertext C_i maps to candidate position (2,2) in $\mathcal{M}_j$, i.e., $(2,2) = \lfloor C_i \rfloor_{2 \cdot \log_2(t)}$. Now, as the current round is j, the mapper computes $\lfloor h(C_i, j) \rfloor_1$. As $\lfloor h(C_i, j) \rfloor_1 = 1$, the mapper stores value "1" in (2,2). Otherwise, he would not touch (2,2).

Input : key-value pairs (i, C_i), PIR values $\{\alpha_1, \ldots, \alpha_t\}$
Output: intermediate key-value pairs (k, v)

for $l := 1$ **to** q **do**
  INITIALIZE $\mathcal{M}_l$;
end
SCANTHROUGHINPUTSPLIT;
foreach *pair* (i, C_i) **do** `//Fill matrices`
  $(\mathcal{X}_i || \mathcal{Y}_i) := \lfloor C_i \rfloor_{2 \cdot \log_2(t)}$;
  for $j := 1$ **to** q **do**
    $bit_j := \lfloor h(C_i, j) \rfloor_1$;
    if $bit_j = 1$ **then**
      $\mathcal{M}_{j, \mathcal{X}_i, \mathcal{Y}_i} := 1$;
    end
  end
end
for $l := 1$ **to** q **do** `//q rounds`
  for $j := 1$ **to** t **do** `//Compute column sums`
    $\sigma_{j,l} = \sum_{\mathcal{M}_{l,1 \le k \le t, j} = 1} \alpha_k$
  end
end
for $j := 1$ **to** t **do** `//Intermediate (k,v) pairs`
  $(k, v) :=$({FILE, j},$\{\sigma_{j,1}, \ldots, \sigma_{j,q}\}$);
  OUTPUT (k, v);
end

Algorithm 2: Computation of PIR matrices $\mathcal{M}$.

Input : reducers' files FILE
Output: decision whether $w \in$ FILE

foreach *file* FILE **do**
  for $i := 1$ **to** q **do**
    if $\lfloor h(C, i) \rfloor_1 = 1$ **then**
      $\mathcal{U}$ reads $s_{\text{FILE}, \mathcal{Y}, i}$;
      $s_i := (s_{\text{FILE}, \mathcal{Y}, i} \cdot b^{-1} \mod p) \mod N$
      `//` $s_i = bit_j$, `see Alg. 2`
      if $s_i = 0$ **then**
        OUTPUT $w \notin$ FILE;
        `//Contradiction`
        break;
      end
    end
  end
  OUTPUT $w \in$ FILE;
end

Algorithm 3: $\mathcal{U}$ decides whether $w \in$ FILE

After all q PIR matrices are filled, the mapper computes for each matrix the t *column sums* $\sigma_{1 \le j \le t, 1 \le l \le q}$ based on $\mathcal{U}$'s input $\{\alpha_1, \ldots, \alpha_t\}$: values α_k with corresponding element $\mathrm{M}_{l,k,j}$ set to "1" are simply added. Finally, the mapper outputs intermediate key-values pairs (k, v). The *key* comprises the name of the file of the InputSplit this mapper was working on, e.g., the file name could be day d, and the number of the column sum of $\mathcal{M}_l$. The *value* consists of a list of the q column sums.

These intermediate key-value pairs will now be input for the reducers during the Reduce phase.

D. Reduce Phase Details

Overview: Recall that there are c InputSplits and therefore c mappers. A single reducer receives from all the c mappers working on the same file all their q column sums for the *same* column. The reducer simply adds these received sums and writes the result into a file which is sent back to $\mathcal{U}$.

Details: For all key-value pairs [({FILE, i},$\{\sigma_{i,1}, \ldots, \sigma_{i,q}\}$)] using the same {FILE, i} as key, the MapReduce framework designates the same reducer. This reducer receives from all c different mappers working on the same file all intermediate key-value pairs with the same key. That is, a reducer receives c pairs which we rewrite as

$$
\begin{gathered}
(\{\text{FILE}, i\}, \{\sigma_{i,1,1}, \ldots, \sigma_{i,q,1}\}) \\
(\{\text{FILE}, i\}, \{\sigma_{i,1,2}, \ldots, \sigma_{i,q,2}\}) \\
\ldots \\
(\{\text{FILE}, i\}, \{\sigma_{i,1,c}, \ldots, \sigma_{i,q,c}\}).
\end{gathered}
$$

Here, for a given $\sigma_{i,j,k}$, i, $1 \le i \le t$, denotes the column, j, $1 \le j \le q$, denotes the round, and k, $1 \le k \le c$, the InputSplit.

Using integer addition, the reducer computes q "final PIR sums" $s_{\text{FILE},i,j} := \sum_{k=1}^{c} \sigma_{i,j,k}, 1 \le j \le q$, and stores values $\{s_{\text{FILE},i,1}, \ldots, s_{\text{FILE},i,q}\}$ into an output file. In conclusion, $s_{\text{FILE},i,j}$ represents the sum of column sums of all the mappers of one particular column i in PIR matrix j. This output file is downloaded by $\mathcal{U}$.

E. Result Analysis

For each outsourced file (day d), user $\mathcal{U}$ retrieves an output file generated by reducers. Now, $\mathcal{U}$ analyzes retrieved files' content to finally conclude which of the outsourced files contain w. Again for ease of understanding, we restrict our description to the analysis of the result generated from PRISM on a single outsourced file called FILE. $\mathcal{U}$ repeats this process with all other results from the other files accordingly.

Definition 8 (Collision)**.** *Assume $\mathcal{U}$ is looking for w, so $C := E_{K_d}(w, 1)$. Similar to hash functions, a* collision *in PIR matrix $\mathcal{M}$ denotes the case of an event where the candidate position $(\mathcal{X}', \mathcal{Y}')$ of another ciphertext $C' \neq C$ matches the candidate position $(\mathcal{X} || \mathcal{Y}) = \lfloor E(w, 1) \rfloor_{2 \cdot \log_2(t)}$ of w in $\mathcal{M}$. That is, $\lfloor C \rfloor_{2 \cdot \log_2(t)} = \lfloor C' \rfloor_{2 \cdot \log_2(t)}$.*

Definition 9 (One-Collision)**.** *A* one-collision *is the event where in an InputSplit a ciphertext $C' \neq E_{K_d}(w,1)$ puts a* 1 *into the same candidate position in $\mathcal{M}$ as $E_{K_d}(w,1)$.*

Overview: The rationale for the result analysis protocol of PRISM is to observe the candidate position of C over q rounds to mitigate the effect of one-collisions. Of particular interest will be rounds where $\lfloor h(C,i) \rfloor_1 = 1$.

First, $\mathcal{U}$ un-blinds all values received from reducers. Based on the result, $\mathcal{U}$ distinguishes two cases.

Case 1.) If a reducer, reducing for a specific file FILE, has returned the value 0 for C's candidate position, then U knows for sure that all mappers have output 0 for this candidate position. Consequently, the candidate position in matrix $\mathcal{M}$ of each mapper is 0. Therefore, C has not been in any of the InputSplits of FILE, and $\mathcal{U}$ reasons $w \notin$ FILE. If C would have been in one InputSplit, then at least the mapper working on this InputSplit would have returned a 1 in this round.

Definition 10 (Contradiction)**.** *Let w be the word $\mathcal{U}$ is looking for, and C its ciphertext. If in some round i, $\lfloor h(C,i) \rfloor_1 = 1$ holds, and the reducer for file* FILE *sends $\mathcal{U}$ a value of* 0 *then this is called a* contradiction*.*

In case of such a contradiction, $\mathcal{U}$ knows for sure that w is not in file FILE.

Case 2.) If, however, this reducer returns a value > 0, then w was in at least one InputSplit *or* a one-collision has occurred in at least one InputSplit. User $\mathcal{U}$ can neither decide $w \notin$ FILE nor $w \in$ FILE with absolute certainty.

U's strategy is to keep the probability for one-collisions low and run multiple rounds q, such that eventually a contradiction occurs ($\Rightarrow \mathcal{U}$ decides $w \notin$ FILE), or, if no contradiction occurs, $\mathcal{U}$ decides $w \in$ FILE with only a small error probability P_{err}.

Details: $\mathcal{U}$ executes Algorithm 3. For each file, $\mathcal{U}$ is only interested in row $\mathcal{Y}$ of matrices $\mathcal{M}$, as only they can refer to candidate position $(\mathcal{X}, \mathcal{Y})$. Therefore, $\mathcal{U}$ keeps only values $\{s_{\text{FILE},\mathcal{Y},1}, \ldots, s_{\text{FILE},\mathcal{Y},q}\}$ and discards the rest. In each round where $\lfloor h(C,i) \rfloor_1 = 1$, un-blinds $s_{\text{FILE},\mathcal{X},i}$ to get value $s_i := \sum_{j=1}^{c} bit_j$. If $s_i = 0$, then we have a contradiction, and $\mathcal{U}$ can infer $w \notin$ FILE. If none of the s_i values has been 0 after all the q rounds, then $\mathcal{U}$ will decide $w \in$ FILE. $\mathcal{U}$ will be wrong with P_{err}.

Note that, although PIR matrices are binary matrices, $\mathcal{U}$ sets $N > c$ to cope with the larger possible values that sums might take due to collisions.

In conclusion, $\mathcal{U}$'s strategy can be summarized by: output $w \notin$ FILE, if $\exists i, s_i = 0$ **or** output $w \in$ FILE, if $\forall i, s_i \neq 0$. We will compute $\mathcal{U}$'s error probability P_{err} for the latter case and dependencies of P_{err} and choices of t and q in Section V-B.

F. Extensions

Saving computation: To save some computation in PRISM, we can modify the hash-based mechanism that determines whether to put a "1" or a "0" in a certain element in $\mathcal{M}$. Recall that the first $2 \cdot log_2(t)$ bit of a ciphertext C are used to determine its position (element) in $\mathcal{M}$. However, instead of computing an expensive hash function $\lfloor h(C_i, j) \rfloor_1$ to get a single bit in round j, we can simply replace the hash and take C's bit on position $(2 \cdot log_2(t) + j)$. Assuming that cipher E has good security properties (each bit of C is "1" with probability $\frac{1}{2}$), this results in the same property as using the hash: eventually two different ciphertexts that collide in $\mathcal{M}$ will differ and lead to a contradiction. We use this computation reduction in our evaluation in Section VI.

Long words: As in related work, e.g., Song et al. [43], we assume for simplicity that words are not longer than our symmetric cipher blocksize. This might not be realistic in some scenarios. To cope with longer words, PRISM could be extended to use (stateful) CBC encryption and to add a hash of each ciphertext to the file. PRISM then searches within hashes and not ciphertexts.

Fuzzy search: PRISM only allows searching for precise words. Using techniques similar to Li et al. [33], i.e., by adding encryptions of variations of words, PRISM could also support *fuzzy* searching.

Counting (finding the number of occurrences)**:** Although PRISM targets primarily search as related work, the use of counters γ can be used for another data mining application: counting. As with a binary search algorithm, $\mathcal{U}$ can launch a set of queries, varying γ, to determine the exact number of occurrences of a word. The number of queries scales logarithmically with the total number of occurrences possible.

V. PRISM ANALYSIS

A. Privacy

We evaluate PRISM's privacy analytically and give short proof sketches. First, we show that PRISM's upload of ciphertexts into the cloud does not reveal any information about plaintexts as in storage privacy. We do this by explaining why our stateful cipher encryption is IND-P1-C0. Second, we show that a set of PRISM queries neither reveals any information about which word $\mathcal{U}$ is looking for, nor whether this word is present (in encrypted form) in the cloud. Hence, we show that PRISM is query private.

1) Storage Privacy:

Lemma 1. *Stateful cipher $E_K(w_i, \gamma_{w_i})$ is IND-P1-C0, where γ_{w_i} is the counter for the current number of occurrence of plaintext w_i, and "," is an unambiguous pairing of inputs.*

Proof (Sketch): During learning, $\mathcal{A}$ submits plaintexts w_i, and encryption oracle $\mathcal{O}_{E_K}$ returns ciphertexts $E_K(w_i, \gamma_{w_i})$. For each submitted plaintext, the stateful cipher increases counter γ_{w_i}. In the challenge phase, $\mathcal{A}$ submits $\{w_0, w_1\}$, and $\mathcal{O}_{E_K}$ returns $E_K(w_b, \gamma_{w_b} + 1)$. However, even if $\mathcal{A}$ can guess $\gamma_{w_b} + 1$, $\mathcal{A}$ does not know secret key K, and $\mathcal{A}$ has never seen $E_K(w_b, \gamma_{w_b} + 1)$ before. As E_K is a good cryptographic encryption function, $\mathcal{A}$ cannot distinguish $E_K(w_b, \gamma_{w_b} + 1)$ from random data.

■

Discussion: The remaining question is how to implement such an unambiguous pairing ",". If the size of the *concatenated* input $|w_i| + |\gamma_{w_i}|$ is less than the cipher's blocksize

(using standard padding), then *concatenation* provides an unambiguous pairing.

We use a stateful cipher construction to achieve IND-P1-C0, because we require encrypted words to be "findable" in a deterministic fashion. While other encryption modes, e.g., CBC or counter mode, *might* also provide IND-P1-C0, they do not allow searching for individual ciphertexts deterministically. As $\mathcal{U}$ cannot know the position of a word in advance, this renders classical encryption modes useless. Similarly, asymmetric probabilistic encryption, e.g., Elgamal, would require additional constructions to pseudo-randomize the ciphertext. Also for performance, we therefore adhere to symmetric encryption.

2) Query Privacy: We transfer the privacy guarantees and constructions of the PIR mechanism of Section III-B to prove query privacy properties. First, note that in PRISM query $Q(w)$ is split into a number of "sub"-queries, namely one query $Q_d(w)$ per file on date d. These queries are processed in parallel by MapReduce, and $\mathcal{U}$ receives output for them separately. Instead of one response R that can be "un-blinded" into a list of files $\mathcal{D}(R) = \{L_1, L_2, \dots\}$ containing w, PRISM provides $\mathcal{U}$ with a set of values $\{s_{\text{FILE},i,j}\}$, one set per file, that are analyzed by $\mathcal{U}$ separately. Therefore, we again restrict our explanations to the case of just one outsourced file queried multiple times.

Lemma 2 explains PRISM's query unlinkability.

Lemma 2. *Given a total of z PRISM-queries $Q(w_i) = \{\alpha_{j,i}\}$ for words w_i, results $R_{\text{File},i} = \{s_{\{\text{FILE},j,k\},i}\}$, ciphertexts $E_K(w_l)$, $z \leq |\mathcal{W}|, 1 \leq i \leq z, 1 \leq j \leq t, 1 \leq k \leq q$, together with PRISM's map and reduce implementation, an adversary $\mathcal{A}$ cannot decide whether $\exists u, v$ such that $w_u = w_v$.*

Proof (Sketch): For each query $Q(w_i)$, values $\{\alpha_{j,i}\}$ are the product of (fresh) random variable b_i and (fresh) random values e_j of the PIR preparation phase. Deciding, whether $\exists u, v$ such that $w_u = w_v$ is equivalent to deciding whether in any of the z sets of α values, there is one value $\alpha_{u,v}$ from the first set and another value $\alpha_{u',v'}$ from the second set, such that either $e_u = e_{u'}$ or $e_{u'} \neq e_{u'}$.

It is computationally infeasible for $\mathcal{A}$ to deduce any information about the mod 2 value of any e_j value. Moreover, e_j values are *randomly* blinded for each individual query, and, therefore, results $\{s_{\{\text{FILE},j,k\},i}\}$, sums of random α, are randomly blinded, too. $\mathcal{A}$ cannot decide whether $e_u = e_{u'}$. ■

Now, we show PRISM's response unlinkability.

Lemma 3. *Given a total of z PRISM-queries $Q(w_i) = \{\alpha_{j,i}\}$ for words w_i, $z \leq |\mathcal{W}|$, results $R_{\text{File},i} = \{s_{\{\text{FILE},j,k\},i}\}$, ciphertexts $E_K(w_l)$, and PRISM's map and reduce implementation, an adversary $\mathcal{A}$ cannot decide for any pair of outputs $(R_{\text{File},i}, R_{\text{File},i'}) = (\{s_{\{\text{FILE},j,k\},i}\}, \{s_{\{\text{FILE},j,k\},i'}\})$ whether $\exists j_1, j_2, k_1, k_2$ such that $s_{\{\text{FILE},j_1,k_1\},i} = s_{\{\text{FILE},j_2,k_2\},i'} \bmod 2$.*

Proof (Sketch): This proof is along the same lines as Lemma 2. As individual α values are indistinguishable from random data, respectively, any pair of $(s_{\{\text{FILE},j_1,k_1\},i}, s_{\{\text{FILE},j_1,k_1\},i'})$ values (sums of α) will be random for $\mathcal{A}$, too. Values $(s_{\{\text{FILE},j_1,k_1\},i}, s_{\{\text{FILE},j_1,k_1\},i'})$ will be different, even if $(s_{\{\text{FILE},j_1,k_1\},i}, s_{\{\text{FILE},j_1,k_1\},i'})$ originate both from queries for the same word w. ■

Finally, we show why queries and results remain confidential.

Lemma 4. *Given a total of z PRISM-queries $Q(w_i) = \{\alpha_{j,i}\}$ for words w_i, results $R_{\text{File},i} = \{s_{\{\text{FILE},j,k\},i}\}$, ciphertexts $E_K(w_l)$, and PRISM's map and reduce implementation, adversary $\mathcal{A}$ can neither discover any word w_i nor any $L_j \in \mathcal{D}(R_{\text{File},i})$ containing w_i.*

Proof (Sketch): As before, we again leverage randomization properties of blinding the α values. Due to blinding, for each query, α and $\{s_{\{\text{FILE},j,k\},i}\}$ values look random to $\mathcal{A}$. Consequently, it is infeasible to discover w_i for any query.

In the case of one outsourced file L_1, the problem of $\mathcal{A}$ learning L_1 can be reduced to deciding whether $\mathcal{D}(R_{L_1,i}) = \emptyset$. However, as $\{s_{\{L_1,j,k\},i}\}$ are random looking sums for each query and response, it is infeasible for $\mathcal{A}$ to decide this problem. As stated above, this holds for the case of multiple files, too. For each file, responses $\{s_{\{\text{FILE},j,k\},i}\}$ values look random to $\mathcal{A}$.

■

B. Statistical Analysis

The remaining question is how $\mathcal{U}$ can choose parameters t and q to get a certain error probability P_{err}. This probability describes the chance that, despite $w \notin \text{FILE}$, $\mathcal{U}$ wrongly outputs $w \in \text{FILE}$ after q rounds without a contradiction, cf., Algorithm 3.

Let n be the number of ciphertexts in one InputSplit, $n := \frac{S_{\text{InputSplit}}}{\text{CipherBlockSize}}$. The total number of ciphertexts stored in the cloud is $(c \cdot n)$.

We consider for simplicity only rounds where $\lfloor h(C,i) \rfloor_1 = 1$, cf., Algorithm 3. If h is a cryptographic hash function, then $\lfloor h(C,i) \rfloor_1 = 1$ in $q' \approx \frac{q}{2}$ of the q total rounds.

First, we recall some basic probabilities. In one InputSplit, the collision probability is $P_{\text{collision}} := \frac{1}{t^2}$.

In one InputSplit, the probability for a one-collision is $P_{\text{one-collision}} := \frac{P_{\text{collision}}}{2}$.

If w is *not* inside a particular InputSplit, the probability that, after inserting the n ciphertexts of that InputSplit into $\mathcal{M}$, the candidate position is *not* set to 1 by collision computes to $P_{\text{InputSplit,no-one-collision}} := (1 - P_{\text{one-collision}})^n$.

Now, if $w \notin \text{FILE}$, i.e., in *none* of the InputSplits, the probability that the candidate position is not set to 1 by collision in *any* InputSplits of FILE is $P_{\text{contradiction}} := (P_{\text{InputSplit,no-one-collision}})^c$.

This is the probability that a contradiction occurs in a single round. If $w \notin \text{FILE}$, the probability that a contradiction occurs in *at least one* of the q rounds is $P_{\text{contradiction,q-rounds}} := 1 - (1 - P_{\text{contradiction}})^q$.

After q rounds without a contradiction, $\mathcal{U}$ automatically decides that w is in FILE. In case that $w \notin \text{FILE}$, and *no* contradiction occurs in q rounds, $\mathcal{U}$ is therefore wrong with

probability $P_{\text{err}} := 1 - P_{\text{contradiction,q-rounds}} = (1 - (1 - \frac{1}{2 \cdot t^2})^{cn})^q$.

Given a certain file size, the size of InputSplits, and the blocksize of the symmetric cipher, $\mathcal{U}$ computes c and n. Therewith, $\mathcal{U}$ can target a false-positive probability by appropriately selecting t and q. We evaluate this using a real-world scenario in Section VI.

VI. Evaluation

To show its real-world feasibility, we have implemented and evaluated PRISM with the scenario described in the introduction. The source code is available for public download [2]. We received 16 days of log data from May 2010 from a small local Internet provider. This provider logs and retains all customers' DNS resolve requests for possible forensic analysis and intrusion detection. Log data is split into files on a daily basis. Each file contains one day of logged 3-tuples: timestamp, customer IP (anonymized by provider for regulatory matters), hostname. The scenario for our evaluation is to use PRISM to upload this data encrypted to MapReduce and perform a search for specific hostnames in a privacy-preserving manner. This is useful for, e.g., "passive DNS analysis" to determine at which day certain command-and-control centers of botnets have been accessed by customer machines, cf., Bilge et al. [5]. The goal of our experiments was to analyze the computational overhead induced by PRISM's privacy mechanism, i.e., the additional time consumed by PRISM over non-privacy-preserving MapReduce.

A. Setup

For the 16 days, the log data contains $\approx 3 \cdot 10^8$ log entries, i.e., $\approx 2 \cdot 10^7$ per file/day. The total space required by all files uploaded into MapReduce using PRISM is 27 GByte, on average 1.7 GByte per file.

Our experiments have been performed on a small "cloud" comprising 1 master computer and 9 slaves. Computers featured a Pentium Dual Core with 2.5 GHz clock frequency and 4 GByte of RAM and running a standard desktop installation of Fedora 11. With this hardware configuration, a total of 18 CPUs were available for maps and reduces. We installed Hadoop version 0.20.2 on our cloud. Being aware that tailoring MapReduce's configuration parameters can have a huge impact on performance, we use the standard, out-of-the-box configuration of Hadoop 0.20.2 without any configuration tweaks. Performance tuning is out of scope of this paper. Similarly, as the InputSplit size is recommended to be between 64 MByte and 128 MByte, we chose $S_{\text{InputSplit}} = 96$ MByte (our InputSplits have to be dividable by $3 \cdot 32$ Byte, since log entries are 3-tuples).

In addition to the evaluation with 96 MByte InputSplits, we also performed a second measurement with slightly larger InputSplits of 120 MByte. We expected a slightly improved performance of PRISM, due to the fact that for the larger files the total number of InputSplits c reduced to less than our 18 available CPUs. Therewith, no additionally costly (re-)scheduling takes places and mappers do not have to process 2 InputSplits sequentially.

TABLE I
Parameters (t, q) to achieve $P_{\text{err}} < 0.01$.

	File size (GByte)							
	0.45	1.21	1.32	1.36	1.38	1.45	1.52	1.67
t	2^{10}	2^{11}						
q	100	60					80	

	File size (GByte)							
	1.78	1.93	2	2.08	2.09	2.14	2.21	2.25
t	2^{12}							
q	20							

Finally, to put timing results into perspective, we implemented and measured a trivial, non-privacy-preserving MapReduce job called *Baseline*. Baseline consists of an empty map phase, where each mapper simply scans over its InputSplit, but does not generate any key-value pairs out of the input. Only at the end of the map phase, a single intermediate key-value pair per mapper is sent to reducers. Upon receipt, reducers discard this key-value pair and write empty files to disk. This trivial baseline only serves in deducing the overhead implied by PRISM, not taking MapReduce specific delays due to rescheduling, speculative execution of backup tasks etc. [20, 40] into account. Note that linear scanning through the entire InputSplit is required in any case, as we assume our data to be generally unordered and unsorted.

For the trapdoor group private information retrieval algorithm, we set $m = 400$ as suggested by Trostle and Parrish [44] for good security and privacy. Our Java implementation is a naive, straightforward implementation using Java's BigInteger without any performance optimizations. As symmetric encryption cipher, we used AES with 256 Bit blocksize from the GNU Crypto Library V2.0.1 [23]. For our IND-P1-C0 extension we reserved $|\gamma| = 2$ Byte and truncated DNS hostnames longer than 30 Byte down to the last 30 Byte.

Simulating $\mathcal{U}$, we computed n and c using blocksize, InputSplit size $S_{\text{InputSplit},}$, and individual file size S_{File}. Assuming that $\mathcal{U}$ targets an error probability of $P_{\text{err}} < 0.01$, we derived t and q. Table I summarizes parameters (t, q) computed for each file individually. Compared to q, we observed that parameter t has a much higher impact on P_{err}, but a comparatively lower impact on computations. Therefore, we increased preferably t than q. Higher values for (t, q) will achieve even smaller values for P_{err}, but Table I shows the computationally "cheapest" combination of (t, q).

B. Results

Computational overhead at the cloud is low as indicated by Figure 2 (PRISM's timing results). We have sorted the 16 files based on their size in an increasing order, i.e, the size of the smallest log file we received from the Internet provider

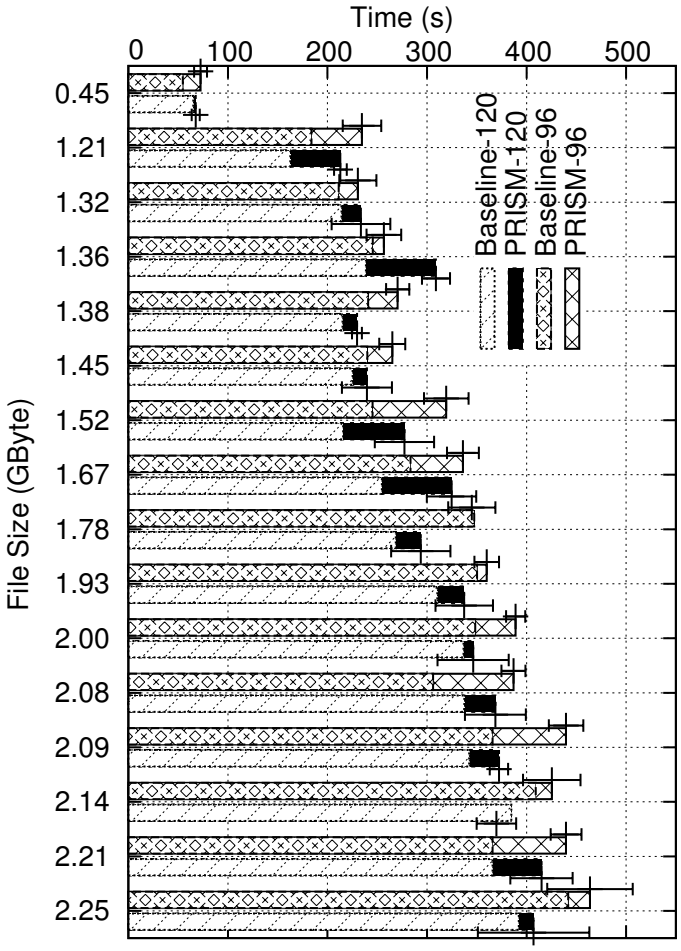


Fig. 2. Wall clock timings for PRISM and Baseline, with $S_{\text{InputSplit}} = 96$ MByte and 120 MByte.

was 0.45 GByte, the largest one was 2.25 GByte. PRISM's execution time was clocked on each file 6 times, respectively, and Fig. 2 shows the average. For each file, Fig. 2 shows two stacked boxes, respectively: the first one for 96 MByte and the second one for 120 MByte InputSplit size. Each of the stacked boxes comprises, first, the baseline timing and, second, the additional time required to run PRISM. To give trust into the evaluation, Fig. 2 also shows 95% confidence intervals drawn right next to each box.

Timings shown in Fig. 2 are "wall clock" timings. This captures the *complete* time elapsed from submitting the PRISM map and reduce classes and starting the job until the end of the reduce phase. In conclusion, the additional overhead over the trivial Baseline MapReduce jobs was on average 11% with a 95% confidence interval of ± 3. The largest overhead seen was 24% over Baseline. This overhead is mostly computational overhead, as there is no difference in disk access between Baseline and PRISM and network volume increases only little by sending slightly larger values during "Reduce". These results do not only show the feasibility of PRISM in practice, but also demonstrate the low overhead implied by PRISM over the non-privacy preserving MapReduce job. We claim that a performance optimized, e.g., not based on Java's BigInteger, implementation, improves performance significantly and furthermore reduces PRISM's overhead.

The simple increase from 96 MByte InputSplit size to 120 MByte InputSplit size has reduced wall clock times for MapReduce jobs on by 9% on average (95% confidence interval of ± 4). Files of size smaller than 2 GByte are split into ≤ 18 InputSplits, and both jobs, PRISM and Baseline, are processed completely in parallel. This indicates that a careful configuration of Hadoop MapReduce's many system parameters, handcrafted and specific to the scenario and jobs to be executed, will lead to substantial performance improvements. This also indicates that in a cloud with more CPUs than in our small setup, the increased number of CPUs will enable to configure way *smaller* InputSplits being processed in parallel. Substantially smaller InputSplits will be beneficial for the overall performance of PRISM or any MapReduce job. However, increasing the number of InputSplits also implies a performance penalty due to (re-)scheduling and coordination activities of the central job tracker, cf., Pavlo et al. [40], so a trade-off has to be found. MapReduce configuration optimizations are, however, out of scope of this paper.

Computational overhead at the User is extremely low in PRISM, too: per file, the preparation of, e.g., $2^{12} = 4096$ α values as required for the underlying PIR scheme is barely measurable (≈ 200 ms) on a PC with 2.5 GHz CPU. During result analysis, $\mathcal{U}$ automatically discards all received values that he is not interested in, i.e., all besides $s_{\text{FILE},\mathcal{Y},1\leq i\leq q}$. For these q values, a total of q Java BigInteger multiplications with modulo are have to be performed. For our examples with $q \leq 100$, this was not measurable at < 1 ms.

Memory consumption for $\mathcal{U}$ is, on the one hand, constant; $\mathcal{U}$ only stores the 256 bit AES key K. On the other hand however, the webserver's memory consumption scales linearly with $O(\mathcal{W})$, i.e., the number of *different* words. This is due to the construction of our stateful cipher that stores counters γ in a hashtable. In our straightforward implementation with Java's standard Hashtable, memory consumption of this hashtable was 548 MByte for the largest log file. While this is certainly a lot of RAM, we conjecture this to be available for a webserver – moreover, as there is a large potential for performance tuning with such data structures.

Communication overhead for PRISM is dominated by the underlying PIR scheme. User $\mathcal{U}$ sends, besides .class files, only the t α values to the cloud. For example, with $t = 2^{12}$ and $m = 400$ Bit, this computes to 200 KByte. The response from the cloud is, for each round, t values of size m. The most expensive configuration in terms of communication in our experiments has been $t = 2^{10}, q = 100$; this results in ≈ 5 MByte communication overhead. Note that communication complexity in the underlying PIR scheme by Trostle and Parrish [44] is linear in the square root of the total table size, i.e., $O(t)$. This can be further reduced by using recursive PIR queries to $O(t^{\epsilon})$, for any $\epsilon > 0$ [32]. Those optimizations as well as amortization techniques discussed by Ostrovsky and Skeith [39] are out of scope.

In conclusion, PRISM is very lightweight for a client using standard PC hardware.

Discussion: While gross timings for queries are high, we point out that most of the time is required by the core MapReduce framework, and PRISM adds only a small computational overhead. Note that on a larger cluster in a more professional environment (hundreds or even thousands of CPUs [28]), all files will be processed in parallel. As you can see in Fig. 2, total time for the 2 GByte file is ≈ 350 s (≈ 6 min). However, already ≈ 340 s are required by MapReduce just to "scan" through the various InputSplits, see Baseline. This has been reported previously, and our results are along the lines of Pavlo et al. [40]. Here, a "grep"-like MapReduce job on 1 TByte of data was clocked at $\approx 1,500$ s on 50 CPUs which would be ≈ 20 times faster than our Baseline. However, Pavlo et al. [40] use a slightly tuned configuration and moreover a more efficient scanning through InputSplits (100 Byte text values instead of 32 Byte binary values in our case) which is known to lead to significant performance increases [30].

VII. Related Work

Cloud Security: Current security research for cloud computing tackles rather traditional issues such as data confidentiality [46], unauthorized access [45] or Denial of Service attacks [27, 34]. Most solutions focus on the protection of the cloud from malicious users. Roy et al. [41] propose a security solution dedicated to MapReduce. Unlike PRISM, the cloud user is the adversary and the cloud is protected against a curious users. In contrast, PRISM considers the *cloud* as a privacy threat.

Private Information Retrieval: Private Information Retrieval (and similarly oblivious transfer and oblivious RAM) has received a lot of attention [9, 15, 16, 22, 25, 32, 38, 39, 39, 42]. In PIR, a user retrieves a specific date from a database. The only "privacy" goal in PIR is access privacy whereby the server should not discover which date a user is interested in. Note that PIR does not ensure privacy of data in the database. PRISM, however, focuses on searching for a word and uses PIR only as a tool.

Searchable Encryption: With searching on encrypted data techniques [6], user privacy is guaranteed thanks to the encryption of the queries and the stored data. However, PRISM offers *higher* privacy guarantees since in existing searchable encryption solutions [4, 6–8, 10, 19, 24, 37, 43], the result ("found" or "not found") originating from a query is known to the adversary. So, the server can obtain some knowledge on users' queries. In PRISM, the adversary does not learn anything about queries or results. Moreover, existing searchable encryption mechanisms *cannot* be extended to leverage from a parallelized cloud setup: while in theory the search on encrypted data itself could be run in parallel on subsets of data, the *combination* of results (as in a reduce phase) is impossible in today's solutions.

VIII. Conclusion

PRISM is the first privacy-preserving search scheme suited for cloud computing. That is, PRISM provides storage and query privacy with very high performance by leveraging parallelism and efficiency of the MapReduce paradigm. Moreover, PRISM is compatible with any MapReduce-based cloud infrastructure since it only requires a standard MapReduce interface (such as Amazon's) without any modification in the underlying system. Thanks to this compatibility, PRISM has been efficiently implemented on an experiemental cloud computing environment using Hadoop MapReduce. Performance of PRISM has been evaluated on that environment through search operations in DNS logs provided by a telco operator. The overhead of PRISM over baseline search operations (without the privacy support) was found to be 11% on the average, acertaining the efficiency of PRISM.

References

[1] Amazon. Elastic mapreduce, 2010. http://aws.amazon.com/elasticmapreduce/.

[2] Anonymized. PRISM source code, 2011.

[3] Apache Hadoop. Welcome to apache hadoop, 2010. http://hadoop.apache.org/.

[4] S.M. Bellovin and W.R. Cheswick. Privacy-enhanced searches using encrypted Bloom filters. Technical Report CUCS-034-07, Department of Computer Science, Columbia University, September 2007. http://tinyurl.com/3v5g9rv.

[5] L. Bilge, E. Kirda, C. Krügel, and M. Balduzzi. Exposure: Finding malicious domains using passive dns analysis. In *Proceedings of 18th Annual Network and Distributed System Security Symposium (NDSS)*, pages 195–211, San Diego, USA, 2011. ISBN 1891562320.

[6] D. Boneh, G. DiCrescenzo, R. Ostrovsky, and G. Persiano. Public key encryption with keyword search. In *Proceedings of Eurocrypt 2004*, pages 506–522, Barcelona, Spain, 2004. LNCS 3027. ISBN 978-3-540-72539-8.

[7] D. Boneh, E. Kushilevitz, and R. Ostrovsky. Publickey encryption that allows pir queries. In *Proceedings of Crypto 2007*, pages 50–67, Santa Barbara, USA, 2007. LNCS 4622. ISBN 978-3-540-74142-8.

[8] G. Brassard, C. Crépeau, and J.-M. Robert. All-or-nothing disclosure of secrets. In *Proceedings on Advances in Cryptology, CRYPTO*, pages 234–238, Santa Barbara, USA, 1986.

[9] C. Cachin, S. Micali, and M. Stadler. Computationally private information retrieval with polylogarithmic communication. In *Proceedings of Advances in Cryptology, EUROCRYPT*, pages 402–414, Prague, Czech Republic, 1999. ISBN ISBN 3-540-65889-0.

[10] Y.-C. Chang and M. Mitzenmacher. Privacy preserving keyword searches on remote encrypted data. In *Proceedings of Applied Cryptography and Network Security (ACNS)*, volume 3531, pages 442–455. LNCS, 2005. ISBN 3-540-26223-7.

[11] Y. Chen and R. Sion. On securing untrusted clouds with cryptography. In *Workshop on Privacy in the Electronic Society*, Chicago, USA, 2010.

[12] Chief Information Officer's Council (CIO). Cloud computing: Benefits and risks of moving federal it into the cloud, 2010. http://tinyurl.com/6hbzx5y.

[13] Chief Information Officer's Council (CIO). Proposed security assessment & authorization for u.s. government cloud computing, 2010. http://tinyurl.com/28yk9fe.

[14] Chief Information Officer's Council (CIO). Privacy recommendations for the use of cloud computing by federal departments and agencies, 2010. http://tinyurl.com/2czrhn3.

[15] B. Chor and M. Naor N. Gilbao. Private information retrieval by keywords. Technical Report TR CS0917, Department of Computer Science, Technion, 1998. http://eprint.iacr.org/1998/003.ps.gz.

[16] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan. Private information retrieval. In *Proceedings of the 36th Annual IEEE Symposium on Foundations of Computer Science*, pages 41–51, Milwaukee, USA, 1995.

[17] Cloud Security Alliance. Security guidance for critical areas of focus in cloud computing, 2009. http://tinyurl.com/ycrchqj.

[18] Cloud Security Alliance. Top threats to cloud computing, 2010. http://tinyurl.com/yer9tvs.

[19] R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky. Searchable symmetric encryption: improved definitions and efficient constructions. In *Proceedings of ACM Conference on Computer and Communications Security, CCS*, pages 79–88, Alexandria, USA, 2006.

[20] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. In *Proceedings of Symposium on Operating System Design and Implementation*, pages 137–150, San Francisco, USA, 2004.

[21] EU. Eu information management instruments, 2010. http://tinyurl.com/3289efv.

[22] Y. Gertner, Y. Ishai, and E. Kushilevitz. Protecting data privacy in private information retrieval. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, pages 151–160, Dallas, USA, 1998. ACM. ISBN 0-89791-962-9.

[23] GNU. The gnu crypto project, 2011. http://www.gnu.org/software/gnu-crypto/.

[24] E.-J. Goh. Secure indexes. Cryptology ePrint Archive Report 2003/216, 2003. http://eprint.iacr.org/2003/216.

[25] O. Goldreich and R. Ostrovsky. Software protection and simulation on oblivious ram. *Journal of the ACM*, 45: 431–473, May 1996. ISSN 0004-5411.

[26] S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984. ISSN 0022-0000.

[27] B. Grobauer and T. Schreck. Towards incident handling in the cloud: Challenges and approaches. In *Proceedngs of the ACM Cloud Computing Security Workshop, CCSW*, Chicago, USA, 2010. ISBN 978-1-4503-0089-6.

[28] Hadoop. Powered by hadoop, list of applications using hadoop mapreduce, 2011. http://wiki.apache.org/hadoop/PoweredBy.

[29] C. Hall, I. Goldberg, and B. Schneier. Reaction attacks against several public-key cryptosystems. In *Proceedings of International Conference on Information and Communication Security*, pages 2–12, Sydney, Australia, 1999. ISBN 3-540-66682-6.

[30] D. Jian, B.C. Ooi, L. Shi, and S. Wu. The performance of mapreduce: An in-depth study. *Proceedings of the VLDB Endowment*, 3(1):472–483, 2010.

[31] J. Katz and M. Yung. Characterization of security notions for probabilistic private-key encryption. *Journal of Cryptology*, 19(1):67–95, 2006. ISSN 0933-2790.

[32] E. Kushilevitz and R. Ostrovsky. Replication is not needed: single database, computationally-private information retrieval. In *Proceedings of 38th Annual Symposium on Foundations of Computer Science*, pages 364–373, Miami Beach, USA, 1997. ISBN 0-8186-8197-7.

[33] J. Li, A. Wang, C. Wang, and W. Lou N. Cao an K. Ren. Fuzzy keyword search over encrypted data in cloud computing. In *Proceedings of INFOCOM Mini-Conference*, pages 441–445, San Diego, USA, 2010. ISBN 978-1-4244-5838-7.

[34] H. Liu. A new form of DOS attack in a cloud and its prevention mechanism. In *Proceedngs of the ACM Cloud Computing Security Workshop, CCSW*, Chicago, USA, 2010. ISBN 978-1-4503-0089-6.

[35] D. McCullagh. Fbi wants records kept of web sites visited, 2010. http://tinyurl.com/yh5omkg.

[36] Official Google Enterprise Blog. Google apps for government, 2010. http://tinyurl.com/3pnabxz.

[37] W. Ogata and K. Kurosawa. Oblivious keyword search. *Journal of Complexity – Special issue on coding and cryptography*, 20:356–371, April 2004. ISSN 0885-064X.

[38] R. Ostrovsky and W. Skeith. Private searching on streaming data. In *Proceedings of Advances in Cryptology, CRYPTO*, pages 223–240, Santa Barbara, USA, 2005. ISBN 3-540-28114-2.

[39] R. Ostrovsky and W.E. Skeith. A survey of single-database private information retrieval: techniques and applications. In *Proceedings of the 10th international conference on Practice and theory in public-key cryptography*, pages 393–411, Beijing, China, 2007. ISBN 978-3-540-71676-1.

[40] A. Pavlo, E. Paulson, A. Rasin, D.J. Abadi, D.J. DeWitt, S. Madden, and M. Stonebraker. A comparison of approaches to large-scale data analysis. In *Proceedings of ACM SIGMOD International Conference on Manage-*

ment of Data (SIGMOD), pages 165–178, Rhode Island, USA, 2009. ISBN 978-1-60558-551-2.

[41] I. Roy, S. Setty, A. Kilzer, V. Shmatikov, and E. Witchel. Airavat: Security and privacy for mapreduce. In *Proceedings of NSDI*, pages 297–312, San Jose, USA, 2010. ISBN 978-931971-73-7.

[42] R. Sion and B. Carbunar. On the computational practicality of private information retrieval. In *Proceedings of Network and Distributed Systems Security Symposium (NDSS)*, pages 1–10, San Diego, USA, 2007.

[43] D.X. Song, D. Wagner, and A. Perrig. Practical techniques for searches on encrypted data. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 44–55, Berkeley, California, 2000.

[44] J. Trostle and A. Parrish. Efficient computationally private information retrieval from anonymity or trapdoor groups. In *Proceedings of 13th International Conference on Information Security*, pages 114–128, Boca Raton, USA, 2010. ISBN 978-3-642-18177-1.

[45] W.Wang, Z. Li, R. Owens, B. Bhargava, and M. Linderman. Secure and efficient access to outsourced data. In *Proceedings of the ACM Cloud Computing Security Workshop, CCSW*, Chicago, USA, 2009. ISBN 978-1-60558-784-4.

[46] A. Yun, C. Shi, and Y. Kim. On protecting integrity and confidentiality of cryptographic file system for outsourced storage. In *Proceedings of the ACM Cloud Computing Security Workshop, CCSW*, Chicago, USA, 2009. ISBN 978-1-60558-784-4.