

Towards Shared Ownership in the Cloud

Hubert Ritzdorf, Claudio Soriente, Ghassan O. Karame, Srdjan Marinovic, Damian Gruber, Srdjan Capkun

Abstract—Cloud storage platforms promise a convenient way for users to share files and engage in collaborations, yet they require all files to have a single owner who unilaterally makes access control decisions. Existing clouds are, thus, agnostic to the notion of shared ownership. This can be a significant limitation in many collaborations because, for example, one owner can delete files and revoke access without consulting the other collaborators.

In this paper, we first formally define a notion of *shared* ownership within a file access control model. We then propose two possible instantiations of our proposed shared ownership model. Our first solution, called *Commune*, relies on secure file dispersal and collusion-resistant secret sharing to ensure that all access grants in the cloud require the support of an agreed threshold of owners. As such, *Commune* can be used in existing clouds without modifications to the platforms. Our second solution, dubbed *Comrade*, leverages the blockchain technology in order to reach consensus on access control decision. Unlike *Commune*, *Comrade* requires that the cloud is able to translate access control decisions that reach consensus in the blockchain into storage access control rules, thus requiring minor modifications to existing clouds. We analyze the security of our proposals and compare/evaluate their performance through implementation integrated with Amazon S3.

Index Terms—Cloud security; Shared ownership; Distributed enforcement; Blockchain technology.



1 INTRODUCTION

Even though the cloud promises a convenient way for users to share files and effortlessly engage in collaborations, it still retains the notion of *individual* file ownership. That is, each file stored in the cloud is owned by a single user, who can *unilaterally* decide whether to grant or deny any access request to that file. However, the individual ownership is not suitable for numerous cloud-based applications and collaborations. Consider a scenario where a number of research organizations and industrial partners want to set up a shared cloud repository to collaborate on a joint research project. If all participants contribute their research efforts to the project, then they may want to share the ownership over the collaboration files so that all access decisions are agreed upon among the owners. There are two main arguments why this may be preferred to individual ownership. First, a sole owner can abuse his rights by unilaterally making access control decisions. The community features a number of anecdotes where users revoke access to shared files from other collaborators. Second, even if owners are willing to elect and trust one of them to make access control decisions, the elected owner may not want to be held accountable for collecting and correctly evaluating other owners' policies. For example, incorrect evaluations may incur negative reputation or financial penalties.

In contrast to individual ownership, we introduce a novel notion of *shared ownership* where n users jointly own a file and each file access request must be granted by a pre-arranged threshold of t owners. We remark that existing cloud platforms, such as

Amazon S3 or Dropbox, provide no support for shared ownership policies, and offer only basic access control lists. In short, they are *agnostic* to the concept of shared ownership. Furthermore, state-of-the-art trust management systems that can support shared ownership policies (e.g., SecPAL [8], KeyNote [11], Delegation Logic [23]) make all access decisions using a *centralized* Policy Decision Point (PDP). This is not suitable for enforcing our shared ownership model, because the user who administrates the PDP can arbitrarily change the policy rules set by the owners and enforce his own policies.

In this paper, we address the problem of *distributed enforcement of shared ownership within cloud storage providers*. By distributed enforcement, we mean enforcement where access to files in a shared repository is granted if and only if t out of n owners separately support the grant decision. Therefore, we introduce the Shared-Ownership file access control Model (SOM) to define our notion of shared ownership, and to formally state the given enforcement problem. We then propose two instantiations of the SOM model to enforce shared ownership policies in a distributed fashion.

This paper extends our previous work [30]. More specifically, we provide additional formal technical details about the SOM model. We also propose a new instantiation of the SOM model, *Comrade*, that leverages functionality from the blockchain in order to reach consensus on access control decisions. Unlike the *Commune* framework proposed in [30], *Comrade* requires cooperation from the cloud provider that is expected to translate access control decisions that reached consensus in the blockchain into storage access control rules. *Comrade*, however, exhibits considerably better performance than *Commune*. We integrate a prototype implementation of *Comrade* within Amazon S3 [1] and compare its performance to the ones of *Commune* [30] with respect to the file size and the number of users. We summarize our contributions as follows:

- *H. Ritzdorf and S. Capkun are affiliated with ETH Zurich, Switzerland. Email: firstname.lastname@inf.ethz.ch*
- *C. Soriente is affiliated with Telefonica Research, Barcelona, Spain. Email: claudio.soriente@telefonica.com*
- *G. Karame is affiliated with NEC Laboratories Europe, Heidelberg, Germany. E-mail: ghassan.karame@neclab.eu*
- *D. Gruber is affiliated with Oneconsult, Switzerland. E-mail: damian.gruber@oneconsult.com*
- *S. Marinovic is affiliated with The Wireless Registry Inc, USA. Email: srdjan@wirelessregistry.com*

- We formalize the notion of shared ownership within a file access control model named SOM, and use it to define a novel access control problem of distributed enforcement of shared ownership in existing clouds.

- We propose a first solution, called **Commune**, which distributively enforces SOM and can be deployed in an agnostic cloud platform. **Commune** ensures that (i) a user cannot read a file from a shared repository unless that user is granted read access by at least t of the owners, and (ii) a user cannot write a file to a shared repository unless that user is granted write access by at least t of the owners.
- We propose a second solution, dubbed **Comrade**, which leverages functionality from the blockchain technology in order to reach consensus on access control decision. **Comrade** improves on the performance of **Commune**, but requires that the cloud is able to translate access control decisions that reached consensus in the blockchain into storage access control rules, thus requiring minor modifications of existing clouds.
- We build prototypes of **Commune** and **Comrade** and evaluate their performance within Amazon S3 with respect to the file size and the number of users.

The remainder of the paper is organized as follows. Section 2 introduces our notion of shared ownership in a file access control model. Section 3 details **Commune** and analyzes its security. In Section 4, we introduce **Comrade** and analyze its provisions. Section 5 evaluates the performance of **Commune** and **Comrade** through an implementation within Amazon S3. In Section 6, we discuss further insights with respect to **Commune** and **Comrade**. Section 7 reviews related work, and we conclude the paper in Section 8.

2 SOM: SHARED-OWNERSHIP FILE ACCESS CONTROL MODEL

In this section, we define the concept of shared ownership, and formally instantiate it in a file access control model dubbed SOM. Our main motivation for constructing this model is three-fold: (i) to precisely define the *ideal* set of features that we believe a model, which enforces shared ownership, should provide; (ii) to formulate the problem of distributed enforcement more precisely by focusing on SOM's formal description; and (iii) to provide a point of reference to scrutinize SOM's enforcement solutions, including our own.

2.1 The Notion of Shared Ownership

In a file system, we see the notion of shared ownership as follows. Each file can have one or more owners, and they collaboratively make an access decision.

To make this notion more precise, let an owner credential denote a pair (O, R) , where R is a tuple $(Subject, File, Action)$, and O is one of $File$'s owners. Intuitively, an owner credential represents a (unilateral) decision by an owner O to grant a request R .

We then define a T -out-of- N file access control policy, also called a *threshold* policy, as follows:

Definition 1 (Threshold Policy). A T -out-of- N (*threshold*) access control policy for a file $File$ is a tuple $(T, Owners, File)$ where T is a number representing a threshold, $Owners$ are the $File$'s owners.

We define an enforcement function $g : Reqs \times TPolicies \times \mathcal{P}(Creds) \mapsto \{grant, deny\}$, where $Reqs$ is a set of requests, $TPolicies$ is a set of threshold policies, and $Creds$ is a set of all

possible credentials. Now we define shared ownership enforcement as follows:

Definition 2 (Shared Ownership Enforcement). An enforcement function g enforces shared ownership over $File$ and its threshold policy $TPolicy$ when $g(R, TPolicy, Creds)$ maps to *grant* iff there are at least T many distinct credentials $(O_1, R), \dots, (O_T, R)$ in $Creds$, where each O_i is in $Owners$ and no two O_i refer to the same owner.

Intuitively, we say that a file access control model enforces shared ownership if it implements a function g that correctly enforces shared ownership.

2.2 SOM's Overview

Given the general notion of shared ownership enforcement from Definition 2, in the following we present a file access control model that adopts this concept in the context of a file access control model. It also further defines how ownership can be delegated and revoked, and how files' thresholds can be changed.

Our model, dubbed SOM, takes files as the only protected resources. We do not focus on directories (or other file groupings). Each file is created by one user with the following request:

$$U \text{ reqs Create}(F)$$

Upon receiving this request, SOM tells a file system to create a file F , assign the user U as the sole owner, and initiate the file's threshold to 1. SOM grants requests for file creation from authenticated users as long as the new file name is unique. To this end, we assume that the file system authenticates U before processing his requests.

SOM allows the ownership over a file to be further shared with, and also revoked from, a user U through the following operations:

- $Delegate(F, U)$ – Delegate ownership of the file F to the user U , i.e., make U one of the owners of F .
- $Revoke(F, U)$ – Revoke ownership of the file F from U , i.e., remove U as an owner of F .

If an owner O wishes to delegate or revoke ownership from U over F , then he issues a credential of the form:

$$O \text{ says Action}(F, U),$$

where $Action$ is either *Delegate* or *Revoke*. Intuitively, one can think of a credential as a certificate by an owner to support an action.

To decide whether a request for an ownership distribution or revocation is in fact enforced for U , SOM consults the file's threshold t to determine how many different credentials U needs from the file's owners. For example, to gain ownership of a file F with $t = 2$, U submits her request:

$$U \text{ reqs Delegate}(F, U)$$

which is granted if two distinct owners of F , for example O and O' , issue the following credentials:

$$O \text{ says Delegate}(F, U)$$

$$O' \text{ says Delegate}(F, U)$$

The full credential and request grammar are defined in Figure 1. To access a file, a user submits the following requests: (i) $Read(F)$

```

credential ::= user says (accessAct | ownsOp | newTOP)
request ::= user reqs (Create(f) | accessOp | ownsOp | newTOP)
accessAct ::= user can accessOp
accessOp ::= Read(f) | Write(f) | Delete(f)
ownsOp ::= Delegate(f, u) | Revoke(f, u)
newTOP ::= NewT(f, t, t)
u ::= String
t ::=  $\mathbb{N}$ 
f ::= String

```

Fig. 1. SOM's credential and request grammar. Words in *italics* are non-terminating symbols.

– Obtain F's content; (ii) Write(F) – Change F's content; (iii) Delete(F) – Delete F. For example, to read F, U submits:

U **reqs** Read(F)

Similarly to granting and revoking ownership, file access requests are granted if t out of N owners issue the corresponding credentials. For example, if the threshold for F is still 2, then U can *read* F, if the following credentials are present:

O **says** U **can** Read(F)
O' **says** U **can** Read(F)

where O and O' are F's owners. We note that in SOM, each of a file's owners can, by default, read that file. However, writing and deleting are still subject to a threshold even for an owner. We find this to be a natural interpretation of shared ownership when compared to unilateral ownership, where an owner has full rights.

Note that successful additions and revocations of the ownership effectively change the number of owners. This, however, does not change the file's threshold. Namely, since adding new owners does not change the threshold t , then the original fraction of owners required to approve file actions is lower. To enable the owners to restore the ratio, or indeed set a new one, the new_T action can be used as follows:

O **says** $\text{New}_T(F, t_{old}, t_{new})$

2.3 Formal Account

Intuitively, we formalize SOM's semantics as follows. We represent a file system state consisting of files, owners and thresholds as a Datalog database [13]. This database consists of a set of relations describing each file's owners and its threshold, and a set of clauses that axiomatize the definitions of shared ownership. We translate a request and credentials into Datalog clauses, which are evaluated over the current state and threshold axioms. For example, file access is granted if a set of credentials supports the grant (expressed as a Datalog query) evaluated over the current state. Facts are added or removed when a set of credentials supports a change of ownership or a change of a particular threshold.

Since SOM's semantics heavily depend on Datalog, we first give a brief overview of Datalog and refer the reader to the more extensive surveys [13]. A Datalog program is a finite set of clauses of the form:

$$S \leftarrow L_1, \dots, L_2,$$

where S and L_i are function-free first-order literals of the form $\text{predicate}(\text{arg}_1, \dots, \text{arg}_n)$. We refer to S as the head of the clause, and to L_i as a body literal. We adopt the following notation: a variable starts with the ? character, a constant starts with a capital letter, and a predicate name starts with a lower-case letter.

A clause with no body literals is called a *fact*. All clauses are safe: all variables that appear in a head literal also appear in at least one body literal. A Datalog program can be split into two sets of clauses: EDB and IDB . EDB is a set of facts whose head literals do not appear as head literals in any other clause. All other clauses are in the IDB set. Intuitively, we think of an EDB as an input for computing all implied facts by the clauses in the IDB set. The declarative semantics of a Datalog program are given by interpreting each clause as a first-order sentence: $\forall \bar{x} L_1 \wedge \dots \wedge L_i \rightarrow S$, and then taking a program to be a conjunction of all its clauses. For each program $\mathcal{P} = IDB \cup EDB$ let $\sigma(IDB, EDB) = \{fact \mid \mathcal{I}(\mathcal{P}) \models fact\}$, where $\mathcal{I}(\mathcal{P})$ represents the first-order translation of \mathcal{P} , and \models is the logical implication.

We formally define SOM's semantics of request evaluations in terms of a labeled transition system (LTS) (S, L, \rightarrow) .

A state $s \in S$ is a tuple $(Files, Users, Owns, Thresholds)$ where $Files$ denotes a set of strings representing file names, $Users$ is a set of users, $Owns$ is a subset of $2^{Users \times Files}$. The $Thresholds$ set is a subset of $2^{Files \times \mathbb{N}}$. For the sake of brevity and presentation, we write $Files_s$, to denote the $Files$ set of the state s (and similarly for other sets of s as well). We can represent a state s as an (EDB) Datalog program s_{EDB} consisting of only the following facts:

$$\begin{aligned}
 \text{file}(F). & \text{ only if } \{F\} \subseteq Files_s \\
 \text{user}(U). & \text{ only if } \{U\} \subseteq Users_s \\
 \text{owns}(U, F). & \text{ only if } \{(U, F)\} \subseteq Owns_s \\
 \text{threshold}(F, N). & \text{ only if } \{(F, N)\} \subseteq Thresholds_s
 \end{aligned}$$

For the sake of simplicity, we assume a fixed set of $Users$ across all states, and we take s_0 to be $(\{\}, Users, \{\}, \{\})$.

A label $e \in L$ is a tuple (R, \mathcal{C}) , where R is a request credential submitted by a user, and \mathcal{C} is a set of available credentials. Credentials can be either submitted by a user, or kept in a separate storage and simply appended to each request.

The \rightarrow set contains all (*valid*) transitions, defined as a triple (s, e, s') . All the necessary and sufficient conditions for valid transitions are given in Figure 2. We note that s and s' are equal in all aspects except if otherwise indicated.

Intuitively, all transition rules (except Create) require that $\sigma(\mathcal{T}(\mathcal{C}) \cup \mathcal{A}[s] \cup s_{EDB}) \models \mathcal{T}(R)$, where s is the state in which the request is received. $\mathcal{T}(C)$ are Datalog clauses generated from e :

$$\begin{aligned}
 \mathcal{T}(U \text{ says } U' \text{ can } \text{accessOp}(F)) &= \text{says}(U, U', \text{accessOp}, F) \\
 \mathcal{T}(U \text{ says } U' \text{ ownsOp}(F, U')) &= \text{says}(U, U', \text{ownsOp}, F) \\
 \mathcal{T}(U \text{ says } \text{New}_T(F, t, t')) &= \text{says}(U, \text{New}_T, F, T, T')
 \end{aligned}$$

In the given translation, accessOp ranges over Write, Read, and Delete; ownsOp ranges over Delegate, and Revoke. The translation of R follows the same idea, except that we do not generate *says* facts but rather queries that should follow from the submitted speech acts.

$$\begin{array}{c}
e = (\text{U reqs } \text{accessOp}(F), \mathcal{C}) \\
\sigma(\mathcal{T}(\mathcal{C}) \cup \mathcal{A}[s] \cup s_{EDB}) \models \mathcal{T}(R) \\
(F \notin \text{Files}_{s'}, \text{ if } \text{accessOp} = \text{Delete}) \\
\hline
s \xrightarrow{e} s' \quad [\text{FAction}]
\end{array}
\quad
\begin{array}{c}
e = (\text{U reqs } \text{New}_T(F, t', t), \mathcal{C}) \\
\sigma(\mathcal{T}(\mathcal{C}) \cup \mathcal{A}[s] \cup s_{EDB}) \models \mathcal{T}(R) \\
(F, t) \in \text{Thresholds}_s, (F, t') \in \text{Thresholds}_{s'}, (F, t) \notin \text{Thresholds}_{s'} \\
\hline
s \xrightarrow{e} s' \quad [\text{NewT}]
\end{array}$$

$$\begin{array}{c}
e = (\text{U reqs } \text{Delegate}(F, U'), \mathcal{C}) \\
\sigma(\mathcal{T}(\mathcal{C}) \cup \mathcal{A}[s] \cup s_{EDB}) \models \mathcal{T}(R) \\
\text{Owns}_{s'} = \text{Owns}_s \cup \{(U', F)\} \\
\hline
s \xrightarrow{e} s' \quad [\text{Delegate}]
\end{array}
\quad
\begin{array}{c}
e = (\text{U reqs } \text{Revoke}(F, U'), \mathcal{C}) \\
\sigma(\mathcal{T}(\mathcal{C}) \cup \mathcal{A}[s] \cup s_{EDB}) \models \mathcal{T}(R) \\
\text{Owns}_{s'} = \text{Owns}_s \setminus \{(U', F)\} \\
\hline
s \xrightarrow{e} s' \quad [\text{Revoke}]
\end{array}
\quad
\begin{array}{c}
e = (\text{U reqs } \text{Create}(F, U), \mathcal{C}) \\
F \notin \text{Files}_s, F \in \text{Files}_{s'}, U \in \text{Users}_s \\
\text{Owns}_{s'} = \text{Owns}_s \cup \{(U, F)\} \\
\text{Thresholds}_{s'} = \text{Thresholds}_s \cup \{(F, 1)\} \\
\hline
s \xrightarrow{e} s' \quad [\text{Create}]
\end{array}$$

where $\text{accessOp} = \{\text{Read}, \text{Write}, \text{Delete}\}$.

Fig. 2. Transition rules for the \rightarrow set of SOM's LTS.

$$\begin{aligned}
\mathcal{T}(\text{U reqs } \text{accessOp}(F) = \text{can}(U, \text{accessOp}, F)) \\
\mathcal{T}(\text{U reqs } \text{ownsOp}(F, U') = \text{ownsOp}(U', F)) \\
\mathcal{T}(\text{U reqs } \text{New}_T(F, t, t') = \text{changeT}(F, T, T')) \\
\mathcal{T}(\text{U reqs } \text{Create}(F)) = \text{create}(U, F)
\end{aligned}$$

$$\begin{aligned}
\text{changeT}(\text{?F}, \text{?T}, \text{?T}_{\text{new}}) \leftarrow \text{file}(\text{?F}), \text{user}(\text{?U}), \\
\text{threshold}(\text{?F}, \text{?T}), [\text{says}(\text{?U}_1, \text{New}_T, \text{?F}, \text{?T}, \text{?T}_{\text{new}}), \\
\dots, \text{says}(\text{?U}_{\text{?T}}, \text{New}_T, \text{?F}, \text{?T}, \text{?T}_{\text{new}}), \\
\text{owns}(\text{?U}_1, \text{?F}), \dots, \text{owns}(\text{?U}_{\text{?T}}, \text{?F}), \\
\text{?U}_1 \neq \text{?U}_2, \dots, \text{?U}_1 \neq \text{?U}_{\text{?T}}, \dots, \text{?U}_{\text{?T}-1} \neq \text{?U}_{\text{?T}}]
\end{aligned}$$

The set $\mathcal{A}[s]$ is a *parameterized* (on s) *IDB* program containing necessary clauses to enforce a *T-out-of-N* access control policy.

The first axiom allows owners to read their files:

$$\text{can}(\text{?U}, \text{Read}, \text{?F}) \leftarrow \text{file}(\text{?F}), \text{owns}(\text{?U}, \text{?F})$$

The second axiom is a template for the *accessOp* operations Read, Write, and Delete:

$$\begin{aligned}
\text{can}(\text{?U}, \text{accessOp}, \text{?F}) \leftarrow \text{file}(\text{?F}), \text{user}(\text{?U}), \\
\text{threshold}(\text{?F}, \text{?T}), \\
[\text{says}(\text{?U}_1, \text{?U}, \text{accessOp}, \text{?F}), \dots, \text{says}(\text{?U}_{\text{?T}}, \text{?U}, \text{accessOp}, \text{?F}), \\
\text{owns}(\text{?U}_1, \text{?F}), \dots, \text{owns}(\text{?U}_{\text{?T}}, \text{?F}), \\
\text{?U}_1 \neq \text{?U}_2, \dots, \text{?U}_1 \neq \text{?U}_{\text{?T}}, \dots, \text{?U}_{\text{?T}-1} \neq \text{?U}_{\text{?T}}]
\end{aligned}$$

Intuitively, this *template* axiom generates the necessary clauses (by substituting *accessOp* with Read, Write, and Delete). The generated clauses are further grounded on $?F$ and $?T$, i.e., on all files and their thresholds. The reason for doing so is to correctly enforce the current (for the given state s) threshold T for a particular file. In other words, we need to generate the correct number of $?U_i$ variables for each file and its threshold in s . To represent this *dynamic* part of a clause (that is dynamically adjusted for each state), we enclose it within $[[$ and $]]$ brackets. We note that the number of variables that need to be generate is given by the $?T$'s value.

The same reasoning applies for the *ownsOp* axioms. We replace $(o|O)\text{wnsOp}$ with $(d|D)\text{elegate}$ and $(r|R)\text{evoke}$, in addition to grounding the clauses on $?F$ and $?T$.

$$\begin{aligned}
\text{ownsOp}(\text{?U}, \text{?F}) \leftarrow \text{file}(\text{?F}), \text{user}(\text{?U}), \text{threshold}(\text{?F}, \text{OwnsOp}, \text{?T}), \\
[\text{says}(\text{?U}_1, \text{?U}, \text{ownsOp}, \text{?F}), \dots, \text{says}(\text{?U}_{\text{?T}}, \text{?U}, \text{ownsOp}, \text{?F}), \\
\text{owns}(\text{?U}_1, \text{?F}), \dots, \text{owns}(\text{?U}_{\text{?T}}, \text{?F}), \\
\text{?U}_1 \neq \text{?U}_2, \dots, \text{?U}_1 \neq \text{?U}_{\text{?T}}, \dots, \text{?U}_{\text{?T}-1} \neq \text{?U}_{\text{?T}}]
\end{aligned}$$

In case of New_T , we ground the following clause on $?F$ and $?T$.

Given these axioms and the transition rules, it follows that SOM represents a correct implementation of an enforcement function g given in Definition 2 for all requests, except when a subject is a file's owner as well and the action is a read action. In this case, an owner is always given access. Clearly, we can easily remove this provision from $\mathcal{A}[s]$, but we argue that it is a natural provision to have in a file access control model.

2.4 Centralized vs. Distributed Enforcement

Given SOM's description, the natural question to consider is how to enforce such a model in a third-party cloud file system that does not endorse shared ownership.

Current state-of-the-art distributed authorization logics—such as SecPAL [8], DKAL [19], Binder [15], KeyNote [11]—that could in principle express SOM's axioms, enforce a policy through a policy decision point (PDP), which evaluates a given set of policies. However, a PDP always has one trusted administrator who has full control over the PDP's policies. This administrator can clearly abuse his powers and modify policies within his PDP and circumvent threshold policies, which defeats the core idea of shared ownership.

We frame this concern as the SOM enforcement problem.

Problem: *How can SOM be enforced without granting one owner unilateral powers?*

3 COMMUNE: DISTRIBUTED ENFORCEMENT OF SHARED OWNERSHIP IN AN AGNOSTIC CLOUD

This section presents Commune, our solution for distributed enforcement of the SOM access control policy in an *agnostic* cloud. As SOM does not specify concrete file access operations, we instantiate Commune with write and read actions. Before introducing our solution, we outline our cloud and attacker model.

3.1 Cloud and Attacker Model

We focus on a cloud storage platform, \mathcal{S} , where a set of users U have personal accounts onto which they upload files. For example,

users might set up their own personal clouds [2], [3], or might create personal accounts in existing public clouds. A user $U \in \mathcal{U}$ can unilaterally decide who has access to files stored on his account. In particular, \mathcal{S} allows each user to define access control policies of the type $p : \mathcal{U} \times \{\text{write}, \text{read}\} \rightarrow \{\text{grant}, \text{deny}\}$. We also assume that \mathcal{S} correctly enforces individual access control policies. This model reflects the functionalities provided by existing cloud platforms, such as Amazon S3.

Since we assume that \mathcal{S} authenticates users, we only focus on internal adversaries. An adversary may try to gain read access to a file even if fewer than t owners have issued the corresponding credentials. We refer to this adversary as a “malicious reader”. Alternatively, an adversary, who has been granted write access by fewer than t owners, may try to publish a file F as if F were authored by a user had been granted write access by t or more owners. We refer to this adversary as a “malicious writer”. We also consider sets of users who collude to escalate their access rights.

3.2 Overview of Commune

Before describing Commune, we make the following observations:

Observation 1. Commune’s files cannot be stored on a single user account.

Following the discussion regarding the centralized enforcement, a single user must not be charged with making unilateral grant and deny decisions. Otherwise, that user may abuse his rights and take unilateral access control decisions. A naïve solution where a file is encrypted (e.g., using a key shared among the owners) and the ciphertext is stored on a single account, allows that account holder to, unilaterally deny read access to the ciphertext. If the ciphertext cannot be read, any mechanism to distribute or recover the encryption key is of no help. We argue, therefore, that Commune cannot use a centralized repository owned by a single user because the repository owner can unilaterally grant or deny access to the files stored therein. Our alternative is to use a “shared repository”, which is an abstraction built on top of the owners’ personal accounts on \mathcal{S} .

Observation 2. Commune cannot support in-place writing.

If Commune were to allow in-place writing, then users who are granted write access could overwrite a file with “garbage”. This would equate to granting users the right to unilaterally delete the file, thus nullifying our efforts to prevent such scenarios. A standard alternative to in-place writing is to introduce “copy-on-write” mechanisms whereby a new file is created upon each file write operation. To optimize performance, Commune implements versioning and splits files into *units* (i.e., the unit of granularity of versioning) so that writing a new version of an existing file, only requires updating the units that have changed with respect to the previous version.

Observation 3. Commune cannot prevent users from disseminating a file through an out-of-band channel.

Access control solutions cannot prevent a user from distributing content through an out-of-band channel. For example, a user who rightfully reads a file can leak it to third parties. Similarly, a malicious writer can write a file and disseminate it through an out-of-band channel. For example, a user can publish files on his account on \mathcal{S} and make them available for others to read. We cannot prevent such behaviour. Commune, however, must at least

allow honest readers, who abide to the protocol specification, to distinguish between the content written by malicious writers and the content written by honest writers.

Given these observations, Commune unfolds as follows. At system setup, users define the set of n owners \mathcal{O} and the threshold t (with $t \leq n$).¹ Commune abstracts the storage space of the owners’ accounts on \mathcal{S} as the “shared repository”. Each owner grants/denies read and write access on his account to users (including other owners) according to his individual access control policy. The distributed enforcement of the SOM access control policy then follows from the enforcement of the individual access policies set by each owner.

To write a file to the shared repository, the writer encodes the file in *tokens* and distributes the tokens to the owners’ accounts. A file is written to the shared repository if and only if the writer successfully distributes the file’s tokens onto at least t owners’ accounts. That is, a user has write access to the shared repository if and only if he has write access to at least t of the owners’ accounts. We refer to such a user as an “authorized writer”.

To read a file from the shared repository, the reader must fetch the file’s tokens from at least t distinct owners’ accounts. Therefore, a user has read access to the file if and only if he has read access to the file’s tokens by at least t owners. We refer to such a user as an “authorized reader”.

To securely enforce shared ownership policies, Commune is designed to fulfil the following properties.

- **P1:** A malicious writer (i.e., a user who has been granted write access by fewer than t owners), must not be able to publish a file F as if F were authored by an authorized writer.
- **P2:** A malicious reader (i.e., a user who has been granted read access to a file F by fewer than t owners), must not be able to recover the file content. This property must also hold in case of *revocation*. Assume that, at the time τ_1 , U has read access to F granted by at least t owners. Also assume that, at the time $\tau_2 > \tau_1$, U has his access rights revoked. This happens if, at the time τ_2 , some of the owners decide to revoke read access to U so that U is left with fewer than t read grants. We must ensure that, starting from time τ_2 , U cannot recover meaningful bits of F . We remark that, as is common for access control systems, we cannot prevent U from storing a local copy of F at the time t_1 and reading it even after his read right has been revoked. Commune must also provide *collusion resistance*. That is, coalitions of users—where no single user is an authorized reader—must not be able to pool their credentials to escalate their read access rights.

Property P1 ensures protection against malicious writers who try to disseminate content despite lacking the required credentials. Property P2 guarantees that malicious readers cannot read content written to the shared repository.

Commune fulfils property P1 by design, through the abstraction of the shared repository and the copy-on-write mechanism (see Section 3.5). Property P2 is fulfilled through two cryptographic building blocks: Secure File Dispersal (SFD), and Collusion Resistant Secret Sharing (CRSS). SFD ensures that malicious readers cannot acquire any information about a file, even if they previously had access to the file and were later revoked. CRSS

1. The selection of owners and the threshold t are outside of our scope. In settings like scientific collaboration, these are agreed upon by the partners.

Algorithm 1 AON-FFT(K, f_1, \dots, f_m)

```

1: Parse  $f_1, \dots, f_m$  as  $f_1^0, \dots, f_m^0$ 
2: for  $r \leftarrow 1$  to  $\log_2 m$  do ▷ round counter
3:   for  $i \leftarrow 0$  to  $\frac{m}{2^r} - 1$  do
4:     for  $j \leftarrow 1$  to  $2^{r-1}$  do
5:        $f_{j+i \cdot 2^r}^r \parallel f_{j+i \cdot 2^r+2^{r-1}}^r \leftarrow E(K, f_{j+i \cdot 2^r}^{r-1}, f_{j+i \cdot 2^r+2^{r-1}}^{r-1})$ 
6:     end for
7:   end for
8: end for
9: return  $f_1^r, \dots, f_m^r$  as  $\bar{f}_1, \dots, \bar{f}_m$ 

```

builds atop SFD and ensures that coalitions of users where no single user has enough credentials to read the file, cannot pool their credentials in order to escalate their read access rights.

In the following, we describe and analyze SFD (Section 3.3) and CRSS (Section 3.4). In Section 3.5, we detail the integration of both building blocks in Commune.

3.3 Secure File Dispersal (SFD)

Information dispersal algorithms [24] encode a file in n chunks so that any t chunks (where $t \leq n$) are sufficient to decode it. However, information dispersal algorithms do not provide any security guarantees if the number of available chunks is smaller than t : any party with fewer than t chunks may still recover meaningful information about the original file's content.

Previous work on securing information dispersal algorithms [25] combines erasure codes with All-Or-Nothing Transforms (AONT) [26]. The latter is an efficient block-wise transformation that maps an n -block bitstring in input to an n' -block bitstring in output (with $n' \geq n$). AONTs are designed in such a way that, unless all the n' output blocks are available, it is hard to recover any of the input blocks.

Existing AONTs [12], [26] leverage block ciphers and rely on the secrecy of a cryptographic key that is embedded within the output blocks. Given all AONT output blocks, the key can be recovered; once the key is known, individual blocks can be reverted, independently of other blocks. Current AONTs, therefore, preserve their all-or-nothing property only for *one time*: knowledge of the cryptographic key allows to revert single output blocks and to recover parts of the original data. This is at odds with our security requirements. As argued before, we cannot prevent users from caching a local copy of the file and reading it at later time when their read rights may have been revoked. However, we still want to provide revocation of a user who only stored the encryption key at the time when he had read access to the file.

We therefore introduce a new scheme, called Secure File Dispersal (SFD), that combines information dispersal algorithms with an AONT that preserves its all-or-nothing property even if the adversary has the encryption key.

Definition. An SFD scheme consists of the following algorithms: $\{c_1, \dots, c_n\} \leftarrow \text{SFD.Encode}(t, n, F, K, \lambda)$. Encodes a file F into n chunks, such that F can be correctly decoded using any t chunks; K denotes a key used in the encoding process and λ is a security parameter.

$F' \leftarrow \text{SFD.Decode}(K, \mathcal{C}, \lambda)$. Takes as input a key K , a set of chunks \mathcal{C} , and security parameter λ ; it outputs a file F' .

Correctness. Given $\{c_1, \dots, c_n\} \leftarrow \text{SFD.Encode}(t, n, F, K, \lambda)$ and $F' \leftarrow \text{SFD.Decode}(K, \mathcal{C}, \lambda)$, we require that if $\mathcal{C} \subseteq \{c_1, \dots, c_n\}$ and $|\mathcal{C}| \geq t$, then $F' = F$.




Fig. 3. Sketch of the AON-FFT scheme where the input consists of $m = 8$ input blocks. Solid circles refer to the block cipher $E(\cdot)$, while empty circles depict its input/output blocks.

Security. We define the advantage of adversary \mathcal{A} as follows:

$$\begin{aligned}
 \text{Adv}_{\text{SFD}}(\mathcal{A}) &= \Pr[f \leftarrow \mathcal{A}(K, \mathcal{C}) \mid K \leftarrow \{0, 1\}^l, l \geq \lambda, \\
 &F = f_1, \dots, f_m \leftarrow \{0, 1\}^{m\lambda}, \\
 &\{c_1, \dots, c_n\} \leftarrow \text{SFD.Encode}(t, n, F, K, \lambda), \\
 &\mathcal{C} \subseteq \{c_1, \dots, c_n\}, |\mathcal{C}| < t, f \subseteq F, |f| \geq \lambda].
 \end{aligned}$$

where $f \subseteq F$ refers to a substring of F . We say that SFD is secure if, for any p.p.t. adversary, its advantage is negligible in the security parameter, i.e., $\text{Adv}_{\text{SFD}}(\mathcal{A}) \leq \text{negl}(\lambda)$. Our security definition captures the scenario where, at an earlier time, \mathcal{A} was given enough chunks to decode F and has cached a copy of the key K , while at current time he is only given fewer than t chunks. Even if \mathcal{A} has the key K , we require the probability that \mathcal{A} recovers any λ consecutive bits of F to be negligible in the security parameter.

Instantiation. Our SFD scheme combines information dispersal techniques with AON-FFT, an all-or-nothing transformation inspired by Fast Fourier Transform.

Let $E : \{0, 1\}^{4\lambda} \rightarrow \{0, 1\}^{2\lambda}$ be a semantically secure block cipher (e.g., $E(\cdot)$ could correspond to 256-bit Rijndael [20], with $\lambda = 128$).² AON-FFT takes as input a symmetric key K (of size 2λ) and m input blocks f_1, \dots, f_m (each of size λ). It executes in $\log_2 m$ rounds and, at each round, applies $E(\cdot)$ to pairs of blocks. Each round is fed with the output of the previous round. The original input f_1, \dots, f_m is treated as the output of round 0; the final output of the algorithm is the output of round $\log_2 m$ (cf. Figure 3). The pseudo-code of AON-FFT is shown in Algorithm 1. We omit the details of the decryption algorithm since it is specular to encryption.

Given the pseudo-code of AON-FFT, our SFD scheme unfolds as follows:

$c_1, \dots, c_n \leftarrow \text{SFD.Encode}(t, n, F, K, \lambda)$. Parse F as f_1, \dots, f_m where each f_i has size λ .

Run $\bar{f}_1, \dots, \bar{f}_m \leftarrow \text{AON-FFT}(K, f_1, \dots, f_m)$. Use the information dispersal encoder to encode $\bar{f}_1, \dots, \bar{f}_m$ in n chunks with reconstruction threshold t .³

$F' \leftarrow \text{SFD.Decode}(K, \mathcal{C}, \lambda)$. Given a set of at least t chunks \mathcal{C} and key K , use the information dispersal decoder to decode blocks $\bar{f}'_1, \dots, \bar{f}'_m$. Run $f'_1, \dots, f'_m \leftarrow \text{AON-FFT}(K, \bar{f}'_1, \dots, \bar{f}'_m)$.

Correctness. If $\{c_1, \dots, c_n\} \leftarrow \text{SFD.Encode}(t, n, F, K, \lambda)$, any subset of at least t chunks $\{c_{i_1}, \dots, c_{i_t}\}$ can be decoded

² The key size is 2λ and the input/output size is also 2λ , totalling 4λ size of input.

³ SFD can leverage any information dispersal algorithm (e.g., Reed-Solomon codes [32]).

into the whole output of AON-FFT, namely $\bar{f}_1 \dots, \bar{f}_m$. Given K , the output of AON-FFT can be decrypted to recover $F = f_1, \dots, f_m$.

Security. Given the construction of our AON-FFT scheme, it is easy to see that each input block depends on all output blocks and on the encryption key. Furthermore, assuming that $E(\cdot)$ is a semantically secure block cipher, for any p.p.t. algorithm \mathcal{A} , we have $\text{Adv}_{\text{SFD}}(\mathcal{A}) \leq \text{negl}(\lambda)$. A full security argument can be found in [30].

Note that a construct similar to AON-FFT, was first mentioned by Rivest [26] and later on used as a ‘‘proof of storage’’ in [31]. Nevertheless, the construction proposed therein can use any pseudo-random permutation in the FFT network. Our AON-FFT requires a keyed permutation, hence a block-cipher. Furthermore, the goal of the adversary in [31] is to recover, in a given amount of time, all *output* blocks. In contrast, the goal of our adversary is to recover any *input* block. This entails a different security analysis.

3.4 Collusion Resistant Secret Sharing (CRSS)

We now introduce our second building block, called Collusion Resistant Secret Sharing (CRSS). Similar to threshold secret-sharing schemes, CRSS allows one party to distribute a secret among a set of designated shareholders, so that any subset of shareholders of size equal to or greater than the threshold can reconstruct the secret. Furthermore, CRSS allows shareholders to issue to other users *delegation* to reconstruct the secret. If a user collects enough (i.e., above the threshold) delegations, he can rightfully reconstruct the secret. However, users cannot pool their delegations to reconstruct the secret, unless one of them has collected enough delegations. In *Commune*, CRSS is used to secret-share the key K used in SFD, in order to achieve collusion resistance.

CRSS is inspired by decentralized Attribute Based Encryption [22] where shares of a secret are *blinded* with shares of 0, such that, if a user collects enough shares for his identity, the blinding cancels out and the secret can be reconstructed.

Definition. Our definition of CRSS builds on top of a *standard* threshold secret-sharing scheme SS with algorithms $\text{SS.Share}(\cdot)$ and $\text{SS.Combine}(\cdot)$, to share and reconstruct a secret, respectively. We assume SS to be secure according to the Game Priv definition by Rogaway et al. [27]. That is, we assume that an adversary has only negligible advantage in identifying which out of two values was (t, n) secret-shared using the $\text{SS.Share}(\cdot)$ algorithm, even if the adversary can corrupt up to $t - 1$ shareholders and access their shares.

CRSS defines the following algorithms:

$\{s_1, \dots, s_n\} \leftarrow \text{CRSS.Share}(s, t, n)$. Shares secret s in a set of n shares $\{s_1, \dots, s_n\}$ with reconstruction threshold t .
 $d_{i,j} \leftarrow \text{CRSS.Delegate}(s_i, U_j)$. Takes as input a share s_i and an user identity U_j . The output is a *delegation* $d_{i,j}$.
 $s' \leftarrow \text{CRSS.Combine}(\{d_{i_1,j}, \dots, d_{i_l,j}\})$. Combines delegations $\{d_{i_1,j}, \dots, d_{i_l,j}\}$ into s' .

Correctness. Given $\{s_1, \dots, s_n\} \leftarrow \text{CRSS.Share}(s, t, n)$ and $s' \leftarrow \text{CRSS.Combine}(\{d_{i_1,j}, \dots, d_{i_l,j}\})$, we require that if $d_{i_p,j} \leftarrow \text{CRSS.Delegate}(s_{i_p}, U_j)$, for $1 \leq p \leq l$ and $l \geq t$, then $s' = s$.

Security. We model the security of CRSS using an adaptation of the Game Priv of [27] and we denote the refined game by Game Priv^* :

Init. The adversary \mathcal{A} submits two messages x_0, x_1 of equal length. The challenger flips an unbiased coin b and runs $\{s_1, \dots, s_n\} \leftarrow \text{CRSS.Share}(x_b, t, n)$.

Find. \mathcal{A} can submit two types of queries. In Type-1 queries, the adversary can corrupt up to $t' \leq t - 1$ shareholders and receives their shares. At this time, \mathcal{A} picks t' indexes $i_1, \dots, i_{t'}$ and receives $\{s_{i_1}, \dots, s_{i_{t'}}\}$. In Type-2 queries, for any fresh identity U_j , the adversary can ask for up to t'' delegations, as long as $t' + t'' \leq t - 1$. \mathcal{A} submits an identity U_j and t'' indexes $i_1, \dots, i_{t''}$, and receives delegations $\{d_{i_1,j}, \dots, d_{i_{t''},j}\}$.

Guess. The adversary outputs his guess b' and wins if $b' = b$.

We define the advantage of the adversary as the probability of its winning minus a half. That is, $\text{Adv}_{\text{CRSS}}^{\text{Priv}^*}(\mathcal{A}) = \text{Prob}[\text{Priv}^* \mathcal{A}] - \frac{1}{2}$. Therefore, we say that CRSS is secure if any p.p.t. algorithm \mathcal{A} has only negligible advantage in winning Game Priv^* .

The above Game Priv^* models a scenario where a set of malicious users, including up to t' shareholders, collects up to t'' delegations for each of their identities. If $t' + t'' \geq t$, the malicious shareholders can produce the missing delegations for any of the colluding user identities, so that the secret can be reconstructed by means of $\text{CRSS.Combine}(\cdot)$. Otherwise, colluding users must not be able to retrieve the secret.

Instantiation. Our CRSS scheme is based on the threshold secret-sharing scheme proposed in [14], which is defined as follows:

$g^x, \{x_1, \dots, x_n\} \leftarrow \text{SS.Share}(-, t, n)$. Pick a cyclic group G of prime order q where the discrete logarithm assumption holds; let $\langle g \rangle = G$. Pick a random $x \in \mathbb{Z}_q$ and set the secret to g^x . Pick a random $t - 1$ -degree polynomial X with coefficients in \mathbb{Z}_q , such that $X(0) = x$. Set the i -th share to $x_i = X(i)$.
 $s' \leftarrow \text{SS.Combine}(\{x_{i_1}, \dots, x_{i_l}\})$. Given l shares $\{x_{i_1}, \dots, x_{i_l}\}$, use polynomial interpolation to recover the secret. That is $s' = g^{\sum_{p=1}^l x_{i_p} \lambda_p}$ where $\lambda_p = \prod_{1 \leq k \leq l, k \neq p} \frac{x_{i_k}}{x_{i_k} - x_{i_p}}$.

Note that in the above scheme, the secret is not given as input to the Share algorithm; rather, it is set to g^x for a randomly chosen x . Given the above algorithms, our CRSS scheme unfolds as follows:

$\{s, s_1, \dots, s_n\} \leftarrow \text{CRSS.Share}(-, t, n)$. Run $\text{SS.Share}(-, t, n)$ to obtain $g^x, \{x_1, \dots, x_n\}$. Pick $H(\cdot) : \{0, 1\}^* \rightarrow G$ to be a cryptographic hash function that maps random strings in G . Pick a random $t - 1$ -degree polynomial Y with coefficients in \mathbb{Z}_q , such that $Y(0) = 0$, and denote $y_i = Y(i)$. The secret is set to $s = g^x$ while each share is set to $s_i = (x_i, y_i)$.
 $d_{i,j} \leftarrow \text{CRSS.Delegate}(s_i, U_j)$. Parse $s_i = (x_i, y_i)$ and output $d_{i,j} = g^{x_i} H(U_j)^{y_i}$.
 $s' \leftarrow \text{CRSS.Combine}(\{d_{i_1,j_1}, \dots, d_{i_l,j_l}\})$. Run $s' \leftarrow \text{SS.Combine}(\{d_{i_1,j_1}, \dots, d_{i_l,j_l}\})$.

Correctness. If $l \geq t$, then $\text{CRSS.Combine}(\{d_{i_1,j}, \dots, d_{i_l,j}\})$ outputs

$$\begin{aligned} s' &= \prod_{p=1}^{p=l} (d_{i_p,j_p})^{\lambda_{i_p}} = \prod_{p=1}^{p=l} (g^{x_{i_p}} H(U_j)^{y_{i_p}})^{\lambda_{i_p}} = \\ &= g^{\sum_{p=1}^{p=l} \lambda_{i_p} x_{i_p}} H(U_j)^{\sum_{p=1}^{p=l} \lambda_{i_p} y_{i_p}} = \\ &= g^k H(U_j)^0 = g^k = s. \end{aligned}$$

Security. The security of CRSS is based on the fact that, in the random oracle model, delegations for different identities cannot be combined to remove the blinding factor from the secret. Assuming that $H(\cdot)$ is modeled as a random oracle and that the discrete logarithm assumption holds in G , we can show that any p.p.t. algorithm \mathcal{A} has only negligible advantage in winning Game Priv^* .

3.5 Commune: Protocol Specification

Recall that **Commune** leverages a shared repository, which is an abstraction of the owners' storage space. The shared repository uses a versioning system so that content cannot be overwritten but only new content can be added. In particular, **Commune** optimizes performance by splitting a file in smaller *units*, and encoding/decoding each unit separately. Therefore, when a new file version is written to the shared repository, the writer only needs to upload the units that have changed from the previous version.

Files written to the repository are encoded in *tokens* and distributed across the owners' accounts. Leveraging the basic ACLs of \mathcal{S} , owners define their individual policy on the tokens in their accounts. The distributed enforcement of the SOM policy is implied by the enforcement of each owner's individual policy on his tokens by \mathcal{S} . Encoding must guarantee both correctness and security of reading operations. Hence, users who are authorized to read at least t tokens must be able to decode the original file; users who are granted read access on fewer than t tokens must not be able to recover its content. Furthermore, users must not be able to pool their credentials to escalate their access rights.

Create a File. File creation requires one user, the file creator, to "bootstrap" the system and write the initial version of the file into the repository. For this reason, we assume that—at the file creation time—the file creator has been granted the right to write new data to each of the owner's accounts on \mathcal{S} .

The file creator splits the file F into k fixed-sized units. For each unit F_i , he runs $\{s_i, s_{i1}, \dots, s_{in}\} \leftarrow \text{CRSS.Share}(-, t, n)$ to produce a fresh secret s_i and n of its shares. Secret s_i is used as a symmetric key to encode the unit F_i in n chunks using SFD. That is, the file creator runs $\{c_{i1}, \dots, c_{in}\} \leftarrow \text{SFD.Encode}(t, n, F_i, s_i, \lambda)$. The token of the unit F_i for the owner O_j is set to (c_{ij}, s_{ij}) (i.e., one chunk outputted by $\text{SFD.Encode}(\cdot)$ and one secret-share outputted by $\text{CRSS.Share}(\cdot)$). Finally, for each owner O_j , the file creator writes $\{(c_{ij}, s_{ij})\}_{i \in [1, \dots, k]}$ to O_j 's account on \mathcal{S} . Each owner, therefore, receives one token for each unit that constitutes F .

Grant/Deny Write Rights. An owner O_j grants write rights to a user U_l by granting to U_l the right to write new data (i.e., tokens) to O_j 's account. Similarly, O_j denies write rights to U_l by denying U_l the right to write new data to O_j 's account.

Update a File. Assume U_l wants to write a new version of a file F . For simplicity, assume that the new version differs from the previous one by only one unit F_i (the case where the old and the new versions differ in several units is handled in a similar fashion). At this point, some owners may allow U_l to write tokens to their accounts while others may not. Let \mathcal{O}^+ be the subset of owners who grant to U_l write rights to their accounts. Similarly, let \mathcal{O}^- be the subset of owners who deny to U_l write rights to their accounts. U_l can, therefore, only distribute tokens to owners in \mathcal{O}^+ . This scenario is equivalent to the case where U_l distributes tokens to all owners in \mathcal{O} , but the ones in \mathcal{O}^- decide to reject the version produced by U_l and make the received tokens unavailable.

U_l is an authorized writer and his version accepted (i.e., considered as written to the shared repository) if and only if $|\mathcal{O}^+| \geq t$. In this case, there are at least t tokens for the new unit, so it may be decoded by users who collect enough credentials. If $|\mathcal{O}^+| < t$, user U_l is not authorized to write and his version is rejected (i.e., considered as not written to the repository), since there are not enough tokens to decode the unit produced by U_l .

Grant/Deny Read Rights. Recall that for each unit F_i , an owner O_j receives the token (c_{ij}, s_{ij}) . O_j can grant to U_l read access to that unit by *endorsing* the token for U_l and granting to U_l read access on the endorsed token. Token endorsement requires O_j to run $d_{ij,l} \leftarrow \text{CRSS.Delegate}(s_{ij}, U_l)$. The endorsed token $(c_{ij}, d_{ij,l})$ is then made available by O_j for U_l to read. If a file consists of multiple units, O_j must endorse all relative tokens for U_l and grant to U_l read access on all endorsed tokens.

O_j can revoke read rights that were previously granted, by denying to U_l the right to read the previously endorsed tokens.

Read a File. If the original file spans several units, U_l must decode each unit separately in order to read the entire file. That is, for each unit, he uses the set of endorsed tokens he can fetch to recover the secret key via $\text{CRSS.Combine}(\cdot)$ and then uses the secret key to decode the unit via $\text{SFD.Decode}(\cdot)$. Note that for an authorized reader to read version x of file F , he must fetch the latest endorsed tokens created up to (and including) version x , for each unit that comprises the file. Assume user U_l is granted read access to $\{(c_{ij_1}, d_{ij_1,l}), \dots, (c_{ij_t}, d_{ij_t,l})\}$. To recover F_i that user runs $s_i \leftarrow \text{CRSS.Combine}(\{d_{ij_1,l}, \dots, d_{ij_t,l}\})$ and then $F_i \leftarrow \text{SFD.Decode}(s_i, \{c_{ij_1}, \dots, c_{ij_t}\}, \lambda)$. U_j proceeds in a similar way to recover all units of F that he has access to.

3.6 Security Analysis

From Sections 3.3 and 3.4, it follows that given t tokens of a file unit F_i (endorsed for a unique user identity), it is possible to recover both the secret key used to encode F_i and its AON-FFT ciphertext, so that the original file can be decrypted. That is, users can read files written by honest writers, if they are granted such right by at least t out of n owners.

Property P1 (cf. Section 3.2) is fulfilled as follows. First, **Commune** uses copy-on-write to prevent writers from overwriting content in the shared repository with garbage. Second, malicious writers (i.e., writers with less than t write permissions) are unable to distribute a file without honest readers detecting it. In other words, a file is considered written if and only if it is correctly encoded in tokens and those tokens are distributed to and endorsed by at least t out of n owners. Any content distributed through other means (e.g., out of band channels) is recognized as malicious by honest readers. We argue that detection of unauthorized files is the only solution for protecting honest readers, because there are no mechanisms to deter malicious writers from disseminating arbitrary content (cf. Observation 3). We also stress that honest readers can easily detect writers that distribute polluted (i.e., non-decodable) tokens. Denial-of-service attacks are, nevertheless, out our scope.

Property P2 is satisfied by combining CRSS and SFD. The former ensures that coalitions of users, where no single user has enough tokens endorsed for his identity, cannot pool their endorsed tokens in order to escalate their access rights. The latter addresses the case where at a time τ_1 a user has access to t or more tokens of a file unit F_i , but at a time $\tau_2 > \tau_1$, his access rights are revoked. That is, at time τ_2 , the user has access to fewer than t endorsed tokens. SFD ensures that even if, at time τ_1 the user may have

cached the key used to encode F_i , he will not be able to decode parts of F_i at time τ_2 . Note that, once a user has access to the file, then he can locally store any plaintext of his choice. Similar to other access control schemes, **Commune** cannot deter this behavior.

Finally, given the guarantees that **Commune** makes for write and read actions, it follows that **Commune** is a (correct) solution for distributed enforcement of the SOM access control policy.

4 COMRADE: BLOCKCHAIN-BASED SHARED OWNERSHIP

In this section, we present an alternative solution for enforcing shared ownership in the cloud by leveraging functionality from the blockchain. Our solution, dubbed **Comrade**, enables a distributed blockchain-based enforcement of the SOM access control policy in a cooperative cloud. Unlike **Commune**, **Comrade** does not assume an agnostic cloud, and requires the cloud operator to cooperate and to interface with the blockchain. Since SOM does not specify concrete file access operations, we instantiate **Comrade** with write and read actions. Before introducing our solution, we provide some background on the blockchain and describe the system model.

4.1 Blockchain and Smart Contracts

The notion of blockchain was originally introduced by the well-known proof-of-work hash-based mechanism that *confirms* cryptocurrency payments in Bitcoin [28]. The PoW-based blockchain ensures that all transactions and their order of execution are available to all blockchain nodes, can be verified by all involved entities and aids the consensus between the parties. Bitcoin’s blockchain fueled innovation, and a number of innovative applications have already been devised by exploiting the secure and distributed provisions of the underlying blockchain. Prominent applications include secure timestamping [6], [7], and smart contracts [16].

Smart contracts refer to binding contracts between two or more parties that are executed by all blockchain nodes. Namely, smart contracts implement state machine replication. Smart contracts typically consist of a self-contained code and persistent storage available to all blockchain nodes. For example, Ethereum [16] is a decentralized platform that enables the execution of arbitrary applications (or contracts) on its blockchain. Owing to its support for a Turing-complete language, Ethereum (which currently also relies on PoW-based consensus) offers an easy means for developers to deploy their distributed applications in the form of smart contracts.

To make smart contracts more powerful, techniques have been developed to securely insert real-world facts into blockchains, such as TownCrier [34]. These facts, such as weather information or flight delays, allow contracts to take real-world events into account and to offer new functionalities.

4.2 Overview of Comrade

In **Comrade**, cloud accounts are not owned by a single user, but by a smart contract that is running within a blockchain. We refer to such a smart contract by *owner contract* and we rely on it to ensure access control as agreed upon by the file owners. The cloud’s PDP makes access control decisions by evaluating a standardized function within the owner contract, as depicted in Figure 4. To grant or deny access rights, the owners submit their votes to the owner contract, which stores them in the blockchain. The PDP’s decision then depends on the access control policy, encoded in the




Fig. 4. Overview of **Comrade**. Access control decisions depend on the evaluation of a smart contract executed within the blockchain.

owner contract, and data stored inside the blockchain, i.e., owners’ votes or securely inserted facts.

To perform an action a on file F in **Comrade**, user U_i proceeds as follows. U_i issues a standard access request to the cloud storage. The request is authenticated using U_i ’s private key. The cloud PDP determines the corresponding owner contract for F and evaluates the `hasAccess()` function inside that owner contract: `hasAccess(F, U_i , a)`.

`hasAccess()` is evaluated based on the contract’s access control policy, the owners’ votes and potentially additional blockchain data. The derived access control decision is then enforced by the cloud’s Policy Enforcement Point (PEP). Notice that the cloud PDP performs this evaluation by locally executing `hasAccess()` on the current state of the blockchain, i.e., the evaluation triggers no action on the blockchain and requires no fees.

The owner contract also manages the users. Users can join the system by sending a request to the owner contract. For every user, the contract’s storage contains the user’s public key, used for authentication and data encryption as explained below. The storage also contains every user’s accounting balance. Finally, the contract contains procedures for initializing and closing the cloud account.

Recall that the owner contract stores the votes inside its storage. To minimize the overhead associated with such a voting scheme (i.e., storage costs in the blockchain), **Comrade** employs a hierarchical file structure and groups files into directories. This allows users to issue a directory-specific vote; votes on directories are valid for all contained files and subdirectories (unless a more precise vote exists). We additionally group users into roles by leveraging role-based access control (RBAC) [18]. RBAC allows full flexibility at higher efficiency as owners only need to vote on access rights for the roles.

Similar to **Commune**, we assume that the cloud provider will enforce access control decisions correctly at all time (although the provider might be interested in learning the contents of files).

Comrade also ensures fair payment by all owners, protect the cloud provider from free-riding, and punishes unfair behaviour. To do so, each user in **Comrade** makes a policy-defined deposit at the owner contract at system setup. The owner contract tracks each user’s balance (e.g., and punishes them for delayed payments). To pay for cloud storage, the owner contract forwards a part of the users’ deposits to a deposit inside the cloud contract.

In turn, the cloud contract deducts the operational costs from the deposit and requests the deposit to be refilled before it reaches zero. Once the deposit reaches zero, access to the cloud resources is denied and after some grace period the cloud resources are released.

Similarly, the owner contract requires users to restock their deposit. Otherwise, the owner contract can impose sanctions, e.g., deny certain access rights or ignore votes in case of an owner. Such sanctions and the payment procedures are defined as part of the owner contract which is visible to all owners at contract creation. Notice here that different accounting policies are feasible. For example, the owners can equally split the costs, can ask users to pay a share of the costs or the policy can dictate usage-based cost sharing where more active users pay more.

In contrast to **Commune**, **Comrade** requires slight changes to the cloud architecture. Namely, the cloud needs to provide a blockchain interface to manage and pay for used cloud resources. To offer such an interface, a single smart contract per cloud provider is sufficient. We refer to such a contract as a *cloud contract*. The cloud can monitor the state of the cloud contract and perform the requested operations. Such an interface seems realistic as cloud providers currently provide more complex interfaces such as command-line tools or web platforms. The cloud also needs a slight modification in its PDP. Access control requests for cloud resources owned by a smart contract are decided by evaluating a function inside the matching owner contract. We refer to this as a blockchain-aware PDP. Overall, **Comrade** only requires minor, inexpensive changes in the cloud infrastructure.

We argue that **Comrade** ensures that the cloud provider cannot be held accountable for collecting and correctly evaluating other owners' policies. For example, incorrect evaluations may incur negative reputation or financial penalties. Instead, all votes are collectively evaluated by the blockchain nodes. Moreover, **Comrade** allows for the first time the implementation of complex, distributed, event-based access control policies that would considerably enrich the cloud offering.

4.3 Comrade: Protocol Specification

We now detail the operations of **Comrade**.

Create a File. During file creation, one user—the file creator—writes the initial version of the file into the repository. This requires the file creator to have write permissions for the directory the file is created in.

The file creator also encrypts the file using a randomly chosen *file key* before uploading it. The encrypted file is uploaded as F using a write action. To securely distribute the file key to U_i , the file creator also uploads wrapped keys F_{k,U_i} containing the file key for file F encrypted with the public key of user U_i . By default a file creator uploads wrapped file keys for all owners. Notice that the access control policy for F_{k,U_i} is defined such that a user U_j can access F_{k,U_i} if and only if $U_j = U_i$ and U_j can access F .

Grant/Deny Write Rights. An owner O_j grants write rights for a resource F (and the associated wrapped file key) to an entity U_l by submitting a corresponding vote to the blockchain. The vote consists of a blockchain transaction $v(O_j, U_l, \text{write}, F)$. Here, F can be a file or directory and U_l can be a single user or a specific role. Similarly, O_j denies write rights for F to U_l by voting against the access. Notice that access to F also implies access to the associated wrapped file key.

Update a File. Assume U_l wants to write a new version of a file F . U_l encrypts F using its file key and issues a write action as described in Section 4.2. In case the owner contract implements a threshold-based access control policy, the request succeeds if there are at least t owner votes in favour.

Grant/Deny Read Rights. Analogously to write rights, an owner O_j grants or denies read rights for a resource F to an entity U_l by submitting a corresponding vote to the owner contract. As mentioned earlier, this vote corresponds to a blockchain transaction $v(O_j, U_l, \text{read}, F)$.

Read a File. Assume U_l wants to read a file F . U_l issues a read request for F and F_{k,U_i} as described in Section 4.2. In case the owner contract implements a threshold-based access control policy, the request succeeds if there are at least t owner votes in favour. U_l decrypts F_{k,U_i} using its private key to obtain the file key and finally decrypts F .

4.4 Security Analysis

We now analyze the security provisions of **Comrade** according to Properties P1 and P2 as defined in Section 3.2.

First, we show that a rational cloud PDP cannot influence the owners' votes. Recall that the owners vote on access control decisions by issuing appropriate blockchain transactions. Such transactions are confirmed in the blockchain by the validators/miners. As required for the security of the underlying blockchain, we assume the standard safety conditions particular to the underlying blockchain technology. For instance, in Proof-of-Work (PoW) based blockchains (e.g., Bitcoin and Ethereum), we assume that the adversary cannot control the majority of the computing power in the network (see [28] for further details). Recall also that the access control decisions are made by the cloud PDP according to the user contract which was previously agreed upon by all owners. These decisions are publicly verifiable and the cloud provider can be held accountable for any diverging decisions.

Property P1 (cf. Section 3.2) is fulfilled as follows. First, similar to **Commune**, **Comrade** uses copy-on-write to prevent writers from overwriting content in the shared repository with garbage. Second, malicious writers (i.e., writers who have been granted write access on a resource by fewer than t owners) are denied access by the cloud PDP following from the execution of the contract.

Property P2 is satisfied by performing access control over the file as an atomic unit. Once a reader loses read access, it cannot read the file anymore as its requests are denied by the cloud PDP as mandated by the owners contract. Even in case the reader had previous access and locally saved the file key, it cannot use this file key to recover meaningful bits of F since it does not have access to the ciphertext. Finally, users cannot collude since the owners' votes are issued to specific users and cannot be reused by other users.

A user anticipating to lose read or write access could try to mount a Denial-of-Service attack against the blockchain to prevent the propagation of the owner vote. This emerges as a challenging task, given the distributed nature/deployment of existing blockchains.

Finally, we note that cloud providers have to evaluate `hasAccess` on every resource access. Since `hasAccess` is Turing-complete, the cloud provider must protect against resource exhaustion attacks, where clients trigger expensive `hasAccess` functions. Therefore, the cloud provider can define a maximal number of execution steps for an evaluation of `hasAccess` and

	New Owner Contract	New Permission	New File
5 Owners	0.59	0.0189	0.00
10 Owners	0.71	0.0189	0.00

TABLE 1
Transaction Fees in USD for our Comrade prototype.

charge the owner contract according to the number of required execution steps. Notice that this is a similar concept as the notion of *gas* in Ethereum [17]. This ensures fair payments across different tenants and defends against resource exhaustion.

5 PROTOTYPE DESIGN & EVALUATION

In this section, we describe prototype implementation of **Commune** and **Comrade** integrated with Amazon S3 [1] and evaluate their performance.

5.1 Commune Implementation

We leverage Amazon S3 to instantiate \mathcal{S} : for each user in \mathcal{U} , we create personal accounts in Amazon S3, into which users can upload content and for which users can define arbitrary access control policies. In our implementation, we use Amazon S3 access control features to distribute tokens from the file creator to the set of owners $\mathcal{O} \subseteq \mathcal{U}$. In particular, we assume that each user sets up (i) one “temporary” folder where other peers are granted write access, and (ii) one “main” folder where endorsed tokens are stored and retrieved. When the file creator wants to distribute a token to owner O_j , he writes the token to O_j ’s temporary folder. Since no other user apart from O_j has read access to the temporary folder, the new token is protected from unauthorized access. At this point, O_j can endorse the token for another user U_l by storing the token in his main folder, and granting read access on it to U_l .

Our **Commune** prototype, implemented in Java, is a multi-threaded client-side interface to repositories hosted on Amazon S3. The client runs on a user’s machine and uploads/downloads content to/from the repositories. The client’s implementation of SFD leverages Rijndael [20] as the underlying block cipher for AON-FFT and systematic Reed-Solomon codes [32] for information dispersal. We chose a symbol size of 16 bytes, and a security parameter $\lambda = 128$ bits.

To optimize performance, our prototype handles file unit operations at a smaller granularity, called *pieces*. During the creation of any file unit, the unit is split into pieces that are processed in parallel. A token for each unit contains one output chunk of SFD for each piece that composes the unit. The piece size w is chosen such that $t\lambda|w$, where λ is the security parameter and t is the required reconstruction threshold. This condition ensures that (i) a piece can be encrypted in an integer number of ciphertext blocks of λ bits, (ii) an encrypted piece can be divided into an integer number of input chunks for the Reed-Solomon encoder, and (iii) the size of each chunk of the Reed-Solomon encoder/decoder is at least λ bits.

5.2 Comrade Implementation

Our python-based implementation of **Comrade** is also integrated with Amazon S3. Here, there is only a single account which is owned by the owner contract. Since Amazon does not support blockchain-aware PDPs yet, we implement the PDP in an Amazon EC2 instance. The PDP has access to the S3 account and makes all

access control decisions based on the current state of the blockchain. Clients vote on the access control policies directly through the blockchain and access resources by contacting the PDP. We rely on Ethereum [16] blockchain running in test mode to avoid paying transaction fees during evaluation. Table 1 summarizes the fees. The creation of our owner contract would cost \$0.59⁴ and \$0.71 for 5 and 10 owners respectively. Granting permissions costs \$0.0189 while uploading a new file incurs no fees.

To support authentication with the PDP, every user registers its public key within the owner contract. Since the storage inside the blockchain is expensive, we use compact elliptic curve cryptography (ECC) (since ECC public keys are smaller than RSA keys). To access a file, a client establishes a TLS connection to the PDP using its registered public key inside a client certificate. The PDP identifies the client based on the key and makes the access control decision by locally evaluating a function of the owner contract. Similar to **Commune**, our **Comrade** prototype breaks units up into pieces.

5.3 Evaluating Single Unit Write/Read

We evaluate the performance of **Commune** and **Comrade** for a single file unit write and read, with respect to (i) the piece size w (default value $w = 128$ bytes), (ii) the reconstruction threshold t (default value $t = 4$), (iii) the number of owners n (default value $n = 10$), and (iv) the size of the file unit $|F_i|$ (default value $|F_i| = 10$ MiB).

We then change one variable at a time to assess its impact on the system performance. For each configuration, we measure the time required (i) to create and upload F_i (denoted by *Write* in our plots), and (ii) to retrieve F_i (denoted by *Read*). These latencies are measured from the initiation of the operation until the output is available either in the repositories (for *Write*) or on a local disk (for *Read*). We control for the effect of caching by uploading random binary streams at each repetition.

During *Read*, the **Commune** client fetches endorsed tokens from t randomly chosen owners. Recall that a (t, n) systematic erasure code outputs t data chunks and $n - t$ parity chunks. Since data chunks need not be decoded, our evaluation accounts for the average-case scenario where the probability that a token contains a data chunk is bounded by $\frac{t}{n}$. Notice that we do not evaluate the time required to grant read rights (i.e., the time required to endorse a token or to submit a blockchain transaction) since it does not depend on any of the considered parameters.

Our results are depicted in Figure 5. For **Commune**, we additionally monitor the runtime of the intermediate steps for a number of configurations as shown in Figure 5(e).

Our evaluation shows that writing a new unit in **Commune** (*Write*) is less expensive than reading it (*Read*) while the order is reversed in **Comrade**. The former effect is due to the overhead of thread synchronization when storing decoded pieces on the local disk while the *Write* performance in **Comrade** is due to the overhead of uploading wrapped keys for all owners.

Impact of the Piece Size: Figure 5(a) shows the impact of the piece size w on the latency. For **Commune**, a smaller w leads to a smaller number of input blocks to the AON-FFT scheme, which results in better performance since AON-FFT requires $\log_2 m$ rounds of encryption for m input blocks. However, we experience higher latencies for very small values of w , especially

4. At time of writing, the Gas price is $2 \cdot 10^{-8}$ ETH where 1 ETH = \$10.54.




Fig. 5. Latency evaluation of our prototype implementations. Each data point is averaged over 20 measurements; where appropriate, we also provide the corresponding 95% confidence intervals. Figure 5(e) splits up the latency into its different components.

in the *Read* operation. This is due to the thread synchronization overhead when writing data to disk. For *Comrade*, a smaller w increases the overhead because of additional synchronization while a larger w cannot benefit from parallelism. Throughout the rest of the evaluation, we set $w = 128$ B for *Commune* and $w = 512$ KiB for *Comrade* since they offer a good performance trade-off as shown in Figure 5(a).

Impact of the Reconstruction Threshold: Figure 5(b) shows the latency impact of the reconstruction threshold t . In *Comrade*, the threshold does not influence latency as shown in Figure 5(b). In *Commune*, the chunk size of the Reed-Solomon encoder increases as t decreases; this results in larger chunk upload and download times. Figure 5(e) also shows that a smaller value of t results in longer encoding and decoding times. On the one hand, during *Write*, small values of t result in larger encoding overhead since the size of the encoding matrix increases. On the other hand, during *Read*, small values of t decrease the probability of recovering data chunks (w.r.t. the probability of recovering parity chunks), which makes decoding slower (cf. Figure 5(e)).

Impact of the Number of Owners: Figure 5(c) shows that latency increases for *Commune*’s *Read* and *Write* as the number of owners grows. The latency increase during *Read* is due to a higher probability of fetching parity codes that take more time to be decoded by the Reed-Solomon decoder. During *Write*, this increase is caused by the creation and distribution of additional tokens from the file creator to the set of owners. The performance of *Comrade* is virtually unaffected by the number of owners since the upload of additional wrapped keys can be performed in parallel—thereby resulting in negligible additional cost.

Impact of the Unit Size: Figure 5(d) shows *Commune*’s and *Comrade*’s latency for different unit sizes. The time required to read/write a unit increases almost linearly with the unit size

Peak Throughput (Mbps)		
	Commune	Comrade
Write	43.39	190.26
Read	29.52	225.37

TABLE 2

Peak throughput. Each data point is the average of 20 measurements.

(Figure 5(d) relies on semi-logarithmic axes). However, the performance of *Comrade* is a magnitude faster than that of *Commune*. The time required to read a 10 MB unit is roughly 4.47 seconds for *Commune*, but only 0.81 seconds for *Comrade*. As shown in Figure 5(e), this stems from the fact that *Commune*’s latency is dominated by the encryption and decryption as part of AON-FFT, while *Comrade* does not leverage AON-FFT and therefore witnesses a considerably lower latency.

5.4 Evaluating Multiple Units Read/Write

We now assess the performance of reading/writing multiple file units in order to determine the peak throughput exhibited by our prototype implementations. Here, we increase the number of concurrently accessed units until the throughput is saturated. We then compute the peak throughput as the maximum aggregated amount of data in bits per second that can be transferred between client and Amazon S3. Table 2 shows that the peak read/write throughput is above 29 Mbps for *Commune* and above 190 Mbps for *Comrade*.

We argue that, while *Commune*’s overhead might be tolerable in low-throughput, high-latency scenarios such as collaborative text editing where users work on content on their local machines (and only periodically synchronize content with the cloud), *Comrade* is a viable option in a wide variety of application scenarios including those with more frequent cloud interactions.

6 DISCUSSION

In this section, we discuss further insights with respect to the design of and possible improvements for **Commune** and **Comrade**.

Transparency to Users: As explained, **Commune** enables users to coordinate access control to cloud content in a distributed manner. We stress that all the operations in **Commune** are implemented by the client application described in Section 5. Users need not “manually” distribute or fetch tokens. In fact, users are only required to set the list of owners for the files they create and to define the access policy on the files for which they are appointed as owners.

In **Comrade**, the owners initially create the owner contract. Afterwards, the owner contract acts as an orchestrator for all users and **Comrade** can act transparently to the users by automatically fetching ciphertexts and the corresponding wrapped keys.

Changing threshold t : To maintain consistency in **Commune**, we do not support the change of threshold t for any file F . If an owner would want to change the threshold, say from t to t' , he would have to compute and distribute new tokens to *all* owners in \mathcal{O} . Then, *all* owners in \mathcal{O} must replace their old tokens with the newly received ones. Since each owner has full rights on its tokens, there is no mechanism to force all owners to accept these changes, and replace their tokens. This can lead to an inconsistent state in which some tokens correspond to a file version with threshold t , while other tokens correspond to another version with threshold t' . Therefore, **Commune** does not support changing the threshold.

In contrast, **Comrade** supports changing threshold t by modifying the owner contract. The owner contract defines the requirements for such a change, e.g., agreement by all owners. Once the requirements for a change are fulfilled, the change takes effect and future evaluations of the owner contract through the cloud PDP use the updated threshold.

Adding/Revoking Owners: Our model assumes that the set of owners \mathcal{O} is defined before file creation. Adding an owner in **Commune** requires that either the original file creator or at least t out of the n owners provide the new owner with his set of tokens. However, revoking ownership rights from an owner, say O_j , may not be feasible since tokens cannot be removed from O_j 's storage on \mathcal{S} without his consent. One possible solution would be to re-encode the file and distribute new tokens to owners in $\mathcal{O} \setminus \{O_j\}$. Nevertheless, similar to the case of changing the threshold t , some of the owners in $\mathcal{O} \setminus \{O_j\}$ may decide to discard the new tokens and keep the old ones—leading to an inconsistent state.

In **Comrade**, owners can be added and revoked through the owner contract. The requirements for adding or revoking owners are mandated by the owner contract which can require e.g., the approval by a majority of owners. Afterwards, the owner list inside the owner contract is updated and new owner votes take effect (or obsolete owner votes are disregarded).

Fine-Grained Per-Version Access Control: **Commune** and **Comrade** enable owners to perform per-version access control. That is, owner O_j can, for example, grant U_l read access to version x of a file F but deny U_l access to F 's version x' . In collaborative scenarios some versions of a given file may contain information only intended for a subset of the users (e.g., due to IPR protection).

Note that, due to versioning, a given unit may span several versions of file F . Nevertheless, this is transparent to the user who only decides whether to grant/deny access to a given version x , while tokens are handled by the client application. In particular

granting/denying read access rights to version x of file F is achieved as follows in both **Commune** and **Comrade**:

To grant read access to version x , the client grants read access to the *most recent version* of each unit that is smaller or equal to x . To deny read access to version x , the client only denies read access to those units that are part of version x but no earlier or later version that U_l should have access to.

7 RELATED WORK

To the best of our knowledge, this is the first work that addresses the problem of distributed enforcement of shared ownership.

Current state-of-the-art access control systems, such as **SecPAL** [8], **KeyNote** [11], and **Delegation Logic** [23], can in principle express t out of n policies. These languages, however, rely on the presence of a centralized PDP component to evaluate their policies. Furthermore, their PDPs cannot be deployed within a third-party cloud platform. As explained in Section 2, these access control systems rely on an administrator to define and manage access control policies. In our setting, this means that a set of owners has to elect one enforcer who has unilateral powers over their files.

Secret sharing schemes [9] allow a dealer to distribute a secret among a number of shareholders, such that only authorized subsets of shareholders can reconstruct the secret. In threshold secret sharing schemes [14], [29], the dealer defines a threshold t and each set of shareholders of cardinality equal to or greater than t is authorized to reconstruct the secret. Secret sharing guarantees security (i.e., the secret cannot be recovered) against a non-authorized subset of shareholders; however, they incur a high computation/storage cost, which makes them impractical for sharing large files.

Information dispersal based on erasure codes [32] are effective tools to enhance the reliability of cloud-based storage systems [4], [5], [21], [33]. Ramp schemes [10] constitute a trade-off between the security guarantees of secret sharing and the efficiency of information dispersal algorithms.

8 CONCLUSION

Even though existing cloud platforms are used as shared repositories, they do not support any notion of shared ownership. We consider this a severe limitation because contributing parties cannot jointly decide how their resources are used.

In this paper, we introduced a novel concept of shared ownership and we described it through a formal access control model, called **SOM**. We then propose two possible instantiations of our proposed shared ownership model. Our first solution, called **Commune**, relies on secure file dispersal and collusion-resistant secret sharing to ensure that all access grants in the cloud require the support of an agreed threshold of owners. As such, **Commune** can be used in existing agnostic clouds without modifications to the platforms. Our second solution, dubbed **Comrade**, leverages the blockchain technology in order to reach consensus on access control decision. Unlike **Commune**, **Comrade** requires that the cloud is able to translate access control decisions that achieved consensus in the blockchain into storage access control rules. **Comrade**, however, shows better performance than **Commune**.

Given the rise of personal clouds (e.g., [2], [3]), we argue that **Commune** and **Comrade** find direct applicability in setting up shared repositories that are distributively managed atop of the various personal clouds owned by users. We therefore hope that our findings motivate further research in this area.

9 ACKNOWLEDGEMENTS

This work was partly supported by the TREDISEC project (G.A. no 644412), funded by the European Union (EU) under the Information and Communication Technologies (ICT) theme of the Horizon 2020 (H2020) research and innovation programme.

REFERENCES

- [1] Amazon Simple Storage Service(Amazon S3). <http://aws.amazon.com/s3/>.
- [2] The Respect Network. <https://www.respectnetwork.com/>.
- [3] WD My Cloud. <http://www.wdc.com/en/products/products.aspx?id=1140>.
- [4] M. Abd-El-Malek, G. R. Ganger, G. R. Goodson, M. K. Reiter, and J. J. Wylie. Fault-Scalable Byzantine Fault-Tolerant Services. In *SOSP*, 2005.
- [5] M. K. Aguilera, R. Janakiraman, and L. Xu. Using Erasure Codes Efficiently for Storage in a Distributed System. In *DSN*, 2005.
- [6] F. Armknecht, J.-M. Bohli, G. O. Karame, Z. Liu, and C. A. Reuter. Outsourced proofs of retrievability. *CCS '14*.
- [7] F. Armknecht, J.-M. Bohli, G. O. Karame, and F. Youssef. Transparent data deduplication in the cloud. *CCS '15*, 2015.
- [8] M. Y. Becker, C. Fournet, and A. D. Gordon. SecPAL: Design and Semantics of a Decentralized Authorization Language. In *Journal of Computer Security (JCS)*, pages 597–643, 2010.
- [9] A. Beimel. Secret-sharing schemes: A survey. In *Third International Workshop on Coding and Cryptology (IWCC)*, pages 11–46, 2011.
- [10] G. R. Blakley and C. Meadows. Security of ramp schemes. In *Advances in Cryptology (CRYPTO)*, pages 242–268, 1984.
- [11] M. Blaze, J. Ioannidis, and A. D. Keromytis. Trust Management for IPsec. In *ACM Transactions on Information and System Security (TISSEC)*, 2002.
- [12] V. Boyko. On the Security Properties of OAEP as an All-or-nothing Transform. In *Proceedings of CRYPTO*, pages 503–518, 1999.
- [13] S. Ceri, G. Gottlob, and L. Tanca. What you always wanted to know about Datalog (and never dared to ask). In *Knowledge and Data Engineering, IEEE Transactions on*, 1989.
- [14] C. Charnes, J. Pieprzyk, and R. Safavi-Naini. Conditionally secure secret sharing schemes with disenrollment capability. In *CCS*, 1994.
- [15] J. DeTreville. Binder, a Logic-based Security Language. In *Proceedings of IEEE Symposium on Security and Privacy*, pages 105 – 113, 2002.
- [16] Ethereum. A Next-Generation Smart Contract and Decentralized Application Platform. Technical report, White Paper, 2016.
- [17] Ethereum community. What is gas? - ethereum homestead 0.1 documentation. <http://ethdocs.org/en/latest/contracts-and-transactions/account-types-gas-and-transactions.html#what-is-gas>.
- [18] D. F. Ferraiolo and D. R. Kuhn. Role-based access controls. In *Proc. of 15th NIST-NSA National Computer Security Conference*, 1992.
- [19] Y. Gurevich and I. Neeman. DKAL: Distributed-Knowledge Authorization Language . In *CSF '08*.
- [20] J. Daemen, and V. Rijmen. AES Proposal: Rijndael. <http://csrc.nist.gov/archive/aes/rijndael/Rijndael-ammended.pdf>.
- [21] J. Kubiawicz, D. Bindel, Y. Chen, S. E. Czerwinski, P. R. Eaton, D. Geels, R. Gummadi, S. C. Rhea, H. Weatherspoon, W. Weimer, C. Wells, and B. Y. Zhao. OceanStore: An Architecture for Global-Scale Persistent Storage. In *ASPLOS*, 2000.
- [22] A. B. Lewko and B. Waters. Decentralizing Attribute-Based Encryption. In *EUROCRYPT*, 2011.
- [23] N. Li, B. N. Groszof, and J. Feigenbaum. Delegation logic: A Logic-based Approach to Distributed Authorization. In *TISSEC*, 2003.
- [24] M. O. Rabin. Efficient Dispersal of Information for Security, Load Balancing, and Fault Tolerance. In *Journal of the Association for Computing Machinery*, 1989.
- [25] J. K. Resch and J. S. Plank. AONT-RS: Blending Security and Performance in Dispersed Storage Systems. In *FAST*, 2011.
- [26] R. L. Rivest. All-or-Nothing Encryption and the Package Transform. In *International Workshop on Fast Software Encryption (FSE)*, 1997.
- [27] P. Rogaway and M. Bellare. Robust computational secret sharing and a unified account of classical secret-sharing goals. In *CCS*, 2007.
- [28] Satoshi Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System.
- [29] A. Shamir. How to Share a Secret? In *Communications of the ACM*, pages 612–613, 1979.
- [30] C. Soriente, G. O. Karame, H. Ritzdorf, S. Marinovic, and S. Capkun. Commune: Shared ownership in an agnostic cloud. *SACMAT '15*, 2015.
- [31] M. van Dijk, A. Juels, A. Oprea, R. L. Rivest, E. Stefanov, and N. Triandopoulos. Hourglass Schemes: how to prove that cloud files are encrypted. In *CCS*, 2012.
- [32] J. H. van Lint. *Introduction to Coding Theory*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1982.
- [33] H. Xia and A. A. Chien. RobuSTore: a Distributed Storage Architecture with Robust and High Performance. In *SC*, 2007.
- [34] F. Zhang, E. Cecchetti, K. Croman, A. Juels, and E. Shi. Town Crier: An Authenticated Data Feed for Smart Contracts. *CCS '16*, 2016.