

SECURE DIGITAL PROVENANCE: CHALLENGES AND A NEW DESIGN

A Thesis

Submitted to the Faculty

of

Purdue University

by

Mohammed M. Rangwala

In Partial Fulfillment of the

Requirements for the Degree

of

Master of Science

May 2014

Purdue University

Indianapolis, Indiana

To my parents, brother, late grandfather,  
family and friends ...  
you keep me going ...

## ACKNOWLEDGMENTS

I pride myself for my achievements during my Computer Science graduate degree at IUPUI. Not only did I have the support of the esteemed faculty members at the university but I also had the opportunity to research on various topics. This has been a long cherished dream. I am grateful to everyone who has backed me during this journey.

First and foremost, I would like to thank my advisor Dr. Xukai Zou for extending his support and guidance throughout my graduate studies. He introduced me to the finer nuances of research and encouraged me to explore the unknown. He challenged my imagination and encouraged me through the difficult and frustrating times. His suggestions and insights into matters related to not only this work but studies in general have been invaluable.

I would like to express my deepest gratitude to Dr. Feng Li and Dr. Rajeev Raje for serving on my advisory committee. Their support and guidance helped me to focus better and make wise decisions in my research and studies.

I want to thank my professors Dr. Mihran Tuceryan, Dr. Arjan Durrezi, Dr. Yao Liang and Dr. Mohammad Al Hasan whose courses have helped me get a better foundation of the subject. I appreciate the cordial smiles and greetings that I had exchanged with the other faculty members in the corridor and at different seminars. A special thanks to Nicole Wittlief, Scott Orr and Nancy Reddington without whose support things would not be running as smoothly.

I was fortunate to receive the prestigious University Fellowship which gave me the impetus to strive harder towards my goals and perform well in all aspects of my degree program.

This journey would not have been possible without the undying support of my parents, brother and family. They have been instrumental in keeping my morale up

and have given me the courage to believe in myself. This has helped me fight against all odds to reach my goals.

Last but not the least, I would like to acknowledge the support of my friends who have always been with me. I could not have done this research without their willingness to stand by me. I have been asked to follow the quote – “To strive, to seek, to find and not to yield”.

## TABLE OF CONTENTS

	Page
LIST OF TABLES . . . . .	vii
LIST OF FIGURES . . . . .	viii
SYMBOLS . . . . .	ix
ABBREVIATIONS . . . . .	x
ABSTRACT . . . . .	xi
1 INTRODUCTION . . . . .	1
1.1 What is Digital Provenance? . . . . .	1
1.2 Main Contributions of this thesis . . . . .	5
1.3 Organization of this thesis . . . . .	6
2 PRELIMINARIES FOR A SECURE PROVENANCE SCHEME . . . . .	7
2.1 Entities involved . . . . .	7
2.2 Properties of a provenance scheme . . . . .	8
2.3 Attack Model . . . . .	11
3 BACKGROUND & RELATED WORKS . . . . .	12
3.1 The Onion scheme . . . . .	13
3.1.1 Overview of the scheme . . . . .	13
3.1.2 Problems with the scheme . . . . .	14
3.1.3 Scheme related to the Onion scheme . . . . .	15
3.2 The PKLC Scheme . . . . .	15
3.2.1 Overview of the scheme . . . . .	15
3.2.2 Problems with the scheme . . . . .	16
4 MUTUAL AGREEMENT SIGNATURE SCHEME . . . . .	18
4.1 Assumptions of our scheme . . . . .	18
4.2 Structure of provenance record and chain . . . . .	18
4.3 Properties satisfied by our scheme . . . . .	24
4.4 Example scenario . . . . .	27
4.5 Auditing activity . . . . .	30
4.6 Advantages of our scheme . . . . .	34
5 SECURITY ANALYSIS . . . . .	38
6 IMPLEMENTATION & EVALUATION . . . . .	43

	Page
6.1 Provenance Record Construction . . . . .	44
6.1.1 Time to create a record and overhead over other schemes . .	44
6.1.2 Overhead of individual operations of the schemes . . . . .	46
6.2 Provenance Chain Auditing . . . . .	48
6.2.1 Concurrent auditing over the chain . . . . .	48
6.2.2 Sequential auditing over the chain . . . . .	49
6.3 Argument of security over performance . . . . .	52
7 CONCLUSION & FUTURE WORK . . . . .	54
7.1 Applications scenarios for our scheme . . . . .	54
7.2 Future Work . . . . .	55
LIST OF REFERENCES . . . . .	57

## LIST OF TABLES

Table	Page
4.1 Summary of advantages of the Mutual Agreement Signature scheme over the Onion and PKLC schemes . . . . .	36

## LIST OF FIGURES

Figure	Page
4.1 An illustration of the structure of provenance . . . . .	21
4.2 Flowchart of a user's actions for creating a provenance record . . . . .	23
4.3 An example of document passing . . . . .	27
4.4 Flowchart of a auditor's actions for auditing a provenance chain . . . . .	32
6.1 Time to create a provenance record in all three schemes . . . . .	45
6.2 Percentage overhead of Mutual Agreement Signature scheme over the Onion and PKLC schemes . . . . .	45
6.3 Percentage of time for each operation in record creation for the Mutual Agreement Signature scheme . . . . .	47
6.4 Percentage of time for each operation in record creation for the PKLC scheme . . . . .	47
6.5 Percentage of time for each operation in record creation for Onion scheme	48
6.6 Provenance chain verification time in all three schemes for file size 50Kb (using concurrent verification) . . . . .	50
6.7 Provenance chain verification time in all three schemes for file sizes 5Mb (using concurrent verification) . . . . .	50
6.8 Provenance chain verification time in all three schemes for file size 50Kb (using sequential verification) . . . . .	51
6.9 Provenance chain verification time in all three schemes for file sizes 5Mb (using sequential verification) . . . . .	51



## SYMBOLS

$K_A$	Public key of user $A$
$K_A^-$	Private key of user $A$
$e_A()$	Encryption using key $K_A$
$h(D_i)$	Cryptographic hash of the document $D_i$
$sign_A()$	Cryptographic signature using key $K_A^-$ . Signature involves the hash of the data, and is short for $sign_i(h(\dots))$
$U_i$	User $i$
$Pr_i$	Provenance record corresponding to user $i$
$Pr_1 Pr_2 \dots Pr_n$	Chain of provenance records
$O_i$	Operations performed by user $i$
$C_i$	Signature of user $i$ over fields of record $Pr_i$
$public_i$	Public key certificate of the user $i$
$I_i$	Keying material for record $Pr_i$
$ChainInfo_i$	Sequence of public keys of users involved in the provenance chain till record $Pr_i$
$S_i$	Symmetric keys for record $Pr_i$
$PubKey_i$	Public keys of previous and next users for record $Pr_i$
$P_i$	Previous signature field representing the previous user in the chain for record $Pr_i$
$N_i$	Next signature field representing the next user in the chain for record $Pr_i$
$IV$	Initialization Vector

## ABBREVIATIONS

PKLC Public-Key Linked Chain

DAG Directed Acyclic Graph

## ABSTRACT

Rangwala, Mohammed M. M.S., Purdue University, May 2014. Secure Digital Provenance: Challenges and a New Design. Major Professor: Xukai Zou.

Derived from the field of art curation, digital provenance is an unforgeable record of a digital object's chain of successive custody and sequence of operations performed on the object. It plays an important role in accessing the trustworthiness of the object, verifying its reliability and conducting audit trails of its lineage. Digital provenance forms an immutable directed acyclic graph (DAG) structure. Since history of an object cannot be changed, once a provenance chain has been created it must be protected in order to guarantee its reliability. Provenance can face attacks against the integrity of records and the confidentiality of user information, making security an important trait required for digital provenance. The digital object and its associated provenance can have different security requirements, and this makes the security of provenance different from that of traditional data.

Research on digital provenance has primarily focused on provenance generation, storage and management frameworks in different fields. Security of digital provenance has also gained attention in recent years, particularly as more and more data is migrated in cloud environments which are distributed and are not under the complete control of data owners. However, there still lacks a viable secure digital provenance scheme which can provide comprehensive security for digital provenance, particularly for generic and dynamic ones. In this work, we address two important aspects of secure digital provenance that have not been investigated thoroughly in existing works: 1) capturing the DAG structure of provenance and 2) supporting dynamic information sharing. We propose a scheme that uses signature-based mutual agreements between successive users to *clearly delineate the transition of responsibility of*

*the digital object* as it is passed along the chain of users. In addition to preserving the properties of confidentiality, immutability and availability for a digital provenance chain, it supports the representation of DAG structures of provenance. Our scheme supports dynamic information sharing scenarios where the sequence of users who have custody of the document is not predetermined. Security analysis and empirical results indicate that our scheme improves the security of the typical secure provenance schemes with comparable performance.

## 1 INTRODUCTION

In this chapter, we give a brief introduction about digital provenance and our work in this thesis.

### 1.1 What is Digital Provenance?

Provenance refers to the origin or earliest known history information of an object. The concept of provenance originates from the field of art and archiving, where it refers to information about the artifact's creation, the chain of custody and modifications performed on it. It has been an important concept in many fields other than art, like science and computing where it is used to trace an object to its origin. It is used in work-flow management systems and processes in physics, astronomy, biology, chemical sciences, earth sciences for maintaining context information, auditing and data replication [1]. It finds applications for intelligent re-use of experiments, fault detection, protection against illegitimate intellectual property claims, detecting plagiarism and identity fraud, and assessments of data quality [2]. Depending on the application domain, different properties of the object can be tracked such as owner information, purpose of its creation, processes undergone, state of the object or material at each stage, etc. Since provenance maintains information about the present and past of an object, it is suitable to assess the object's trustworthiness [3].

Digital provenance is the provenance associated with digital objects which can be resources in hardware, software, documents, databases and other entities. It maintains information about the chain of successive custody of the object with different users and the sequence of operations performed on it. It can store functional data such as the process results as well as non-functional data such as the performance of each

step. Most computer systems track information for error correction and debugging, as discussed in [4], such as:

- i) Operating systems store logs of important system events which help in system administration and intrusion detection.
- ii) File systems store information about file creation, modifications performed, permissions to the file, etc.
- iii) Version Control Systems record information about the modifications made to different objects.
- iv) Web browsers store history information about the web pages visited and when.

These can be considered as different forms of provenance information, which are application specific. However, each of these systems does not provide a definition for provenance.

Digital Provenance finds applications in a number of areas [5]. Some of them are:

- i) Verification of scientific data and experiments [6–14];
- ii) Supporting or facilitating data sharing [15–18];
- iii) Copyright clearance [19];
- iv) Legal proceedings involving data [20];
- v) Tracking operations on data in cloud environments [21–23];
- vi) Recently in facilitating data mining [24]; tracing system activities in Android devices [25]; stream management [26, 27];

Provenance systems are specially designed application-specific frameworks used to collect, analyze and store all metadata information of an object. They can then be queried to obtain the history information, perform audits and validation checks and detect faults. Research in provenance has focused on developing such frameworks for

a variety of systems like work-flow management, grid computing, file systems, cloud systems. We mention some of these in Chapter 3.

Digital provenance introduces some challenges with respect to its definition, management and security [4]. Some of these challenges are:

- i) **Completeness:** Since provenance contains history information, it is necessary to define how much recorded information is considered to be enough. Depending on the application, it maybe necessary to record the output of each individual operation performed on the object. This is important because completeness of the provenance will define the complexity of a provenance system.
- ii) **Reliability:** Provenance must be reliable since it finds applications in fault detection, identity theft and plagiarism detection, etc. It is necessary for the provenance to be secure against any kind of tampering after it has been created.
- iii) **Heterogeneity:** A digital object can undergo several operations that may produce different meta-data information. It may also be recorded at different levels of granularity. Thus, provenance can contain heterogeneous information which introduces the challenge of uniform consistent representation.
- iv) **Portability:** Since provenance is associated with a digital object, it must be bound to it. As the object moves in the system, its provenance must move along with it. This requires the provenance to support portability in the system.
- v) **Dynamic nature:** Different users may operate on the digital object at different points in its lifetime. In a distributed information network or wireless sensor network, the sequence of nodes through which the object passes is predetermined (to a certain extent). But in different scenarios, it may be possible that the sequence is dynamic. The structure of the provenance and the provenance system itself must be able to support this.

Provenance can be represented as a causality graph that connects different objects with edges that describe the process by which the object transformation took place

[28]. This forms an immutable directed acyclic graph (DAG) structure. Although an object keeps changing with operations performed on it, its history information does not and so, provenance is immutable. The DAG structure is justified since an object can be copied to multiple instances (or provided as input to multiple processes) and it can be created from a combination of objects (or from outputs of multiple processes). In a graph, these cases represent a node having multiple children or multiple parents respectively. Since history information does not repeat, the graph does not have any cycles. There is no established standard for representing provenance information, but XML is most popularly used [1]. The existing limited security mechanisms for provenance do not appropriately apply to DAG structures.

Depending on the application domain, provenance can be more or less sensitive than the data object itself. For e.g., in an employee review system, the sequence of managers who have added to the review must not be disclosed to the employee. Thus, the ownership information in the provenance chain in such a scenario must be kept confidential. Here, the provenance is more sensitive than the document it is associated with. Consider another example of a professor's recommendation for a student's university application. The recommendation document itself needs to be kept confidential from the student, but the provenance containing the information of the professor(s) can be disclosed. In such a scenario, the document is more sensitive than its provenance. Apart from this, like other information security subjects, digital provenance requires integrity and availability, along with suitable and efficient representation. In this respect, the security requirements of provenance differ from those of traditional data [28]. Thus, a general scheme for secure provenance is needed, which can be modified depending on the application scenario.

Recent research in provenance has focused on developing provenance generation, storage and management frameworks in different fields but limited work has focused on the security and privacy issues related to it. We recognize that two aspects have not gained enough attention: 1) capturing the DAG structure of provenance and 2) supporting dynamic information sharing. In this work, we propose a scheme that uses



signature-based mutual agreements between successive users to secure the provenance chain. It is an interactive protocol that clearly delineates the transition of responsibility of the digital object as it is passed along the chain of users. A related provenance scheme was proposed by Hasan et. al. [29]. This scheme is referred to by Wang et. al. [30] as the *Onion* scheme due to its layered provenance format. They showed that the Onion scheme has certain weaknesses and proposed a linked chain structure of provenance using public keys. This scheme is referred to as the Public-Key Linked Chain (*PKLC*) scheme [30]. The PKLC scheme works well for distributed information systems but cannot handle all the properties required in other digital systems. Our solution extends their work and solves the problems associated with it.

## 1.2 Main Contributions of this thesis

The contributions of our work can be summarized as:

- A signature-based mutual agreement scheme is proposed to form links between provenance records. Our scheme provides better security than the Onion scheme [29], and, is an extension and improvement over the PKLC [30] scheme to provide secure provenance in digital systems other than distributed information networks.
- Our scheme can adequately support the representation of DAG structures of provenance.
- It can also support dynamic information sharing scenarios where the next user to whom the data will be passed is not predetermined. A summary of the advantages of our scheme is provided in Table 4.1.
- An analysis of the security of our scheme is provided to show that it satisfies the security requirements of a provenance scheme.
- Experimental evaluations are provided for the overhead of our scheme. The overhead of our scheme for provenance record creation is a little more than the

other schemes but we argue that it can be outweighed by the security provided by our scheme. The results show that it performs better than the Onion scheme for provenance chain verification.

### 1.3 Organization of this thesis

The rest of this thesis is organized as follows. In Chapter 2, we discuss the fundamental concepts involved in a provenance scheme, the important properties required for its security and an attack model that must be considered. We highlight the previous proposed mechanisms in Chapter 3. Chapter 4 describes our mutual agreement signature scheme, along with an example and discusses its properties. We analyze the security of our scheme with respect to the attack model in Chapter 5. Performance evaluation and comparison with existing schemes is provided in Chapter 6. Chapter 7 talks about future work and concludes the thesis.

## 2 PRELIMINARIES FOR A SECURE PROVENANCE SCHEME

In this chapter we describe and provide definitions for some fundamental concepts related to digital provenance. We discuss the different entities involved in a secure provenance scheme; the security properties required from it; and a general attack model for building a secure provenance scheme.

### 2.1 Entities involved

A **document**  $D$  is a data item such as a file, database tuple or network packet for which provenance is to be generated and maintained. In this work we use the term document abstractly; its exact form is domain and application-specific.

**Provenance** of a digital document is an account of all the actions performed on it right from the point of creation. Each access to the document can create a provenance record  $Pr$ ; multiple such records are maintained in order as a **provenance chain**  $Pr_1|Pr_2|\dots|Pr_n$ . Provenance of a document forms a directed acyclic graph (DAG) structure [28]. We refer to the provenance of the document as a ‘chain’ in this work because records are arranged sequentially, but they may not be linearly linked to each other. A provenance record stores in it an account of the operations performed by a user on the document, and relevant information that help maintain links between the different records that are part of the chain.

**Users** are the entities who have or have had custody of the document. They may perform operations on the documents, e.g. create, rename, read, write, delete in the case of a file system. The user who first creates the document and is associated with the first record of the provenance chain is referred to as the **owner** of the provenance. This is different from the current owner of the document. In this work we refer to owner as the owner of the provenance chain.

An **auditor** is an entity who can check all provenance information to verify the lineage of a document. An auditor performs an **auditing** activity, which involves traversing the provenance records in the chain and checking their fields to ensure that the chain has not been tampered with. Different users may trust different auditors with sensitive information, thus, a document has a set of auditors who can access different sensitive fields in the records.

**Outsiders** are entities who do not have access to the documents, and subsequently should not have access to any part of its provenance.

An **adversary** has access to the provenance and wants to alter it in some way for malicious intents but remain undetected. An adversary may be a user who has already contributed to the provenance chain or an outsider.

## 2.2 Properties of a provenance scheme

After discussing the fundamental entities, we discuss the properties that a scheme must provide for the provenance data. Groth et. al. [31] identified a set of properties that any provenance system must provide. We list them here:

- i) **Verifiability:** The provenance scheme must be able to verify a process with respect to the users involved, operations performed and results obtained.
- ii) **Accountability:** The scheme must hold the user accountable for his/her actions, i.e. a user should not be able to repudiate any actions.
- iii) **Reproducibility:** The provenance should contain enough information for it to be possible to reproduce the same results if the sequence of operations recorded is re-executed.
- iv) **Preservation:** Since provenance contains history information, it must be maintained for a sufficiently long period of time.
- v) **Scalability:** For large scale applications, a large amount of provenance data may be generated which requires the scheme to be scalable.

- vi) **Generality:** It is possible for a wide variety of meta-data to be generated from an application and the provenance scheme should be general enough to be able to capture them.
- vii) **Customizability:** The scheme must allow customization to be able to record any application-specific details at different levels of granularity.
- viii) **Portability:** Provenance is associated with a digital object, and there must be a mechanism to ensure that they cannot be separated. Along with this, the scheme must allow the provenance to move in the system when the data moves.

We now discuss the properties that a scheme must provide for securing the provenance chain which are more related to our work. These are extended from the fundamental general properties of data security. We mention the properties here to get an understanding of the security required for a provenance chain and discuss how our scheme achieves them in Section 4.3.

**Confidentiality:** A provenance record may contain sensitive information regarding the operations performed on the document as well as its ownership history that should not be revealed to unauthorized entities. The sensitivity of these fields is domain and application-specific. For e.g., in an employee review system, the sequence of managers who have added to the review must not be disclosed to the employee. Thus, the ownership information in the provenance chain in such a scenario must be kept confidential. This is different from the confidentiality of the document itself. Thus, provenance and the document may have different confidentiality requirements. The properties that are required are:

- i) An auditor should be able to verify the complete lineage of the document, without access to the sensitive information in the records.
- ii) Since different users may trust different auditors, the sensitive information may not be revealed to all auditors.

**Integrity:** Since provenance contains history information which is immutable, integrity of the provenance is the most important property that a scheme must satisfy. There are three types of integrity associated with provenance [30]:

- i) Immutability (Chain Integrity): The provenance chain once formed should not be modifiable, i.e. the order of the records cannot be changed.
- ii) Data Integrity: The information in the individual provenance records should not be tampered with.
- iii) Non-repudiation (Origin Integrity): A user's action in the chain cannot be undone, i.e. the user cannot repudiate his actions.

**Availability:** Provenance is associated with a document and when it is passed between users, the provenance chain is passed along with it. The scheme must ensure that when the document is passed between users, the chain remains intact and is not modified without being detected in the auditing activity.

**Efficiency:** Depending on the application domain, the provenance generation process and scheme participates either when operations are being performed on the document (when outputs of individual operations must be recorded) or after all operations have been performed. In both cases, the provenance scheme adds a computational overhead on the application. The scheme must be designed such that the overhead is not significant.

As seen in this discussion, the representation and properties of provenance can be different from those of the document it is associated with. Each individual provenance record can have different confidentiality requirements, whereas the integrity and availability of all records in the chain must be protected as the chain grows. This introduces new challenges in the security of provenance, making it different from the security of traditional data.

### 2.3 Attack Model

Here we briefly discuss some of the goals of the adversary in a digital provenance scheme similar to discussions in [29, 30]. A detailed analysis of attacks and their prevention in our scheme is discussed in Chapter 5. A provenance scheme must consider an adversary with the intentions of,

- i) obtaining confidential information from the provenance records about the operations performed on the document;
- ii) obtaining information about the ownership history of the document;
- iii) using fake or stolen key-pairs to make their own provenance records un-verifiable;
- iv) modifying existing records (tampering or changing order of records) or adding forged information to the existing provenance chain;
- v) selectively removing a certain part of the preceding provenance chain.

Our scheme should be designed such that these goals are either prevented or made detectable to an auditor in the auditing activity.

These preliminaries lay the foundation for understanding the existing work done in secure digital provenance (Chapter 3) as well as our scheme (Chapter 4).

### 3 BACKGROUND & RELATED WORKS

Before discussing our scheme in detail, we discuss some of the research that has been conducted with respect to the security of digital provenance and give an insight into the motivation for our work.

Research has been done to develop conceptual frameworks and models for provenance management [11, 13, 14, 31–42]; to identify the security requirements of provenance systems [4, 5, 43] and provenance management and data forensics in cloud environments [21–23, 44–47].

Hasan et. al. [20] were among the first to propose the concept of secure provenance. Although provenance had been studied in many applications and fields, they identified that the security issues had not been considered. They defined the properties required from a secure provenance scheme along with a threat model and challenges. Braun et. al. [28, 48] discussed some of the essential characteristics of provenance and how it is different from other data in terms of security. They were among the first to recognize that the provenance graph is a directed acyclic graph (DAG) structure to which traditional security measures cannot be directly applied.

Kairos [2] is an architecture for securing the data authorship and temporal information in provenance records suited for work-flow-based grid computing environments. It uses techniques from public key infrastructure (PKI) such as certificate authorities, digital signatures and time stamping protocols to protect provenance records. Kairos has a centralized architecture involving a certificate authority and a time stamp authority, which in combination are responsible for time stamping and signing a provenance record for user of the grid application. The architecture aims at protecting the provenance record but does not give details about the structure of the record itself.



Sultana et. al. [49, 50] proposed a lightweight method for detecting provenance forgery in wireless sensor networks. Since such systems are power and memory constrained it is necessary for the provenance management to be efficient in storage and transmission and not be computation intensive. In this scheme, they use bloom filters to encode provenance information to be able to detect packet drop attacks. Sultana et. al. also proposed a method to securely transmit provenance information in streaming media [27]. Though these works are not directly related, they consider some of the characteristics of confidentiality and integrity preservation required by our scheme.

Alharbi et. al. [51] proposed a privacy-preserving data provenance scheme to ensure the security of provenance for documents on remote servers. The main focus of the scheme is to preserve the privacy of the users through the use of hash chains and group signature techniques, and employs the use of a trusted authority and trusted servers. Our scheme makes use of only a trusted auditor but is not focused at remote document operations.

### 3.1 The Onion scheme

The Onion scheme [29, 52] is closely related to our work. In this section, we give a brief overview of this scheme and discuss its shortcomings.

#### 3.1.1 Overview of the scheme

The Onion scheme was the first to define a concrete structure of a provenance record. Each individual provenance record in a chain of records has the following structure:

$$Pr_i = \langle U_i, O_i, h(D_i), C_i, public_i, I_i \rangle$$

We limit our discussion of the fields of the records in this chapter, since they will be elaborated in Chapter 4 when we discuss our scheme. Here,  $U$  contains the user's

information,  $O$  is a representation of the operations performed on the document,  $h(D)$  is the hash of the document,  $C$  is the checksum of the record,  $public$  stores a public key certificate of the user and  $I$  contains keying material. Multiple such records arranged sequentially form the provenance chain. The checksum field also contains the checksum of the previous record in the chain and this makes it an incremental chained signature mechanism. Each record uses the checksum signature over the previous record's signature to preserve the integrity of the complete chain. Since the checksum field is layered, it gives the chain an onion-like structure [30]. The  $U$  and  $O$  fields may contain sensitive information. They can be kept confidential using symmetric keys with auditors, which are stored in  $I$ .

### 3.1.2 Problems with the scheme

The Onion scheme has certain weaknesses [30]. First, it cannot protect the outermost layers of the provenance chain, i.e. the newest records. An insider attacker can easily extract a prefix of the complete chain, sign over the signature of the last record in the extracted chain and insert a new record. The flaw comes from the fact that this scheme is not based on a hand-off mechanism when the document is passed between users. Consecutive records in the chain are loosely linked to each other.

Second, the scheme requires the trusted auditor(s) to maintain user-key relationship which violates the confidentiality of the users. Our scheme involves a *ChainInfo* field in the provenance record, which sequentially stores the public keys of all users involved in the preceding chain. This field provides the keys necessary to perform the operations with the records, but does not reveal any identity information of the users even to the trusted auditor(s).

Also, the scheme cannot support the DAG structure of provenance. It is restricted to the scenario of a single document being passed along a chain of users. Our scheme overcomes these weaknesses by introducing a mutual agreement mechanism, when the document is passed between users, which creates strong links between their prove-

nance records. It also uses multiple fields for signatures to handle the case of multiple parents or children in a DAG structure.

### 3.1.3 Scheme related to the Onion scheme

Syalim et. al. [53] proposed a scheme based on the Onion scheme. In this scheme, every document and its provenance passed between users is signed by the previous user as well as the owner of the provenance to preserve the integrity of the chain. They define a path-based policy as well as a compartment policy for providing access control on the provenance graph structure. The provenance records are encrypted using multiple keys that are handled by the provenance owner.

The shortcoming of this scheme is that it makes use of the involvement of the provenance owner to satisfy the properties of the system. This heavy involvement of the owner at each step of the provenance scheme is undesirable. This is avoided in our scheme, through a mutual agreement mechanism between only the users who are involved in passing the document at a particular time. Also, a large number of encryption and signature operations are performed which make the scheme inefficient.

## 3.2 The PKLC Scheme

The Public-Key Linked Chain (PKLC) scheme [30] is most closely related to our work. We give a brief discussion of this scheme.

### 3.2.1 Overview of the scheme

The PKLC scheme [30] is based on advancements to the Onion scheme applied to a distributed information network. It uses a record format similar to that of the Onion scheme, and has the following structure:

$$Pr_i = \langle U_i, O_i, h(D_i), S_i, PubKey_i, C_i \rangle$$

Though the fields of the PKLC scheme are originally named differently, here we use the same notations as those of the Onion scheme for convenience. The structure is suited for distributed information networks or wireless sensor networks.  $U$  is similar to that of the Onion scheme where it stores information about the node performing operation on the document.  $O$  contains information of every individual operation that is performed on the document by the same node.  $h(D)$  is the hash of the document and  $S$ , similar to the  $I$  field stores symmetric keys used to encrypt sensitive information.  $C$  contains the signature of the node over the complete record. The scheme differs in the manner in which it links records of the chain. A record can contain the previous and next user's public keys in the *PubKey* field to link the records.

### 3.2.2 Problems with the scheme

The links between the records are weak since they are formed with only the public keys of the users. This scheme is suitable for distributed information or wireless sensor networks where each user initially knows the identity of the next user to which the document passes. In such a scenario, the weak links are enough to preserve the integrity of the chain. The auditors know the path of information flow among users and can thus verify the chain of records. But this cannot be applied directly to dynamic information sharing scenarios where the next user is not predetermined. Our scheme builds on this drawback of the PKLC scheme. It does not require the identity of the next user to be known, but ensures the integrity by having the users engage in a mutual agreement scheme at the moment when the document is passed between them.

To denote the owner of the provenance chain, the first record contains the public key of the owner in the previous field as well. If applied to a general scenario, it is susceptible to an owner forgery attack. An adversary can remove the records of the chain and claim to be the owner by creating a record with his/her own public key in the previous field. Thus, a stronger mechanism is required for representing and

distinguishing the owner of the provenance from other users. Our scheme handles this by involving the auditor.

We have now laid the foundation for a secure provenance scheme and discussed previously proposed provenance schemes along with their shortcomings. We now discuss our proposed design in the next chapter.

## 4 MUTUAL AGREEMENT SIGNATURE SCHEME

In this section we present our scheme for secure digital provenance.

### 4.1 Assumptions of our scheme

Before we discuss our scheme in detail, we mention our assumptions.

- i) Our scheme considers only the format of the provenance records and chain. It does not focus on the storage and maintenance of the chain. Storage systems such as PASS [33], Flogger [35] can be used for this purpose.
- ii) The provenance generation and storing mechanism is not compromised. Our scheme focuses on securing the provenance from attacks after it has been created and stored securely.
- iii) The keys used for signatures and encrypting the fields are never compromised or revoked.
- iv) The document and its provenance are inseparable, i.e. when a document is passed, the provenance chain is also passed with it. This must be maintained by the provenance storing mechanism.
- v) Our scheme relies on transitive trust defined in [30]. That is, pairs of users involved in the document passing trust each other. Thus, we assume that consecutive pairs of users do not collude.

### 4.2 Structure of provenance record and chain

The provenance chain is composed of a sequence of individual provenance records. Each record stores fields that contain information about the user, operations per-

formed, chain of custody of the document thus far, and a representation of the previous and next users in the chain. It has the following structure:

$$Pr_i = \langle U_i, O_i, h(D_i), ChainInfo_i, S_i^*, C_i, P_i^+, N_i^* \rangle$$

Figure 4.1 gives a representation of the structure of the the provenance chain and each individual provenance record for the scheme. Each of the fields of a record is explained as follows.

$U_i$  contains identity information about the user  $i$  who creates this provenance record.

This information is specific to an application domain. For a file system provenance record,  $U_i$  includes user ID, process ID, ipaddress, port, host, time, and so on.

$O_i$  gives a representation of the sequence of operations or modifications performed on the document by user  $i$ . This is also dependent on the application domain. For the file system provenance record,  $O_i$  includes a file diff, log of changes or operations, or any other reversible representation [29]. It can also contain subfields for representing the operations performed by different processes under the same user as in [30].  $O_i$  contains a reversible representation of the operations if the application domain supports it. By reversible we mean that given document  $D_i$  and  $O_i$ , it is possible to obtain  $D_{i-1}$ .

$U_i$  and  $O_i$  contain information about the identity of the user and the operations performed, which may be sensitive to the application. They may be encrypted, in which case the  $S_i$  is used.

$h(D_i)$  is the cryptographic hash of the contents of the document  $D_i$  after user  $i$  performs all operations. A hash function is a one way function that is almost always unique for different documents. As the document is modified along the chain, a hash in each record uniquely represents the state of the document at that instant.

**ChainInfo<sub>i</sub>** is a representation of the chain of custody of the document tracked from its origin. It is a sequence of the public keys of all users involved with this document from the owner of the document  $U_1$  to the current user  $U_i$  represented as  $K_{Aud}|K_1|K_2|\dots|K_i$ .  $K_{Aud}$  is the public key of the auditor, who must begin the provenance chain. The purpose of  $K_{Aud}$  will become clear in the further discussion.

**S<sub>i</sub><sup>\*</sup>** stores symmetric keys that may be used to encrypt the sensitive  $U_i$  and  $O_i$  fields. We adopt the broadcast encryption scheme of [29] to regulate the access for different auditors. Instead of creating multiple encrypted versions of the sensitive fields for each auditor, user  $i$  encrypts them with a symmetric key  $K_s$ , and then stores copies of  $K_s$  encrypted with the keys  $K_{Aud_j}$  of the respective auditors. The \* indicates there may be zero or more symmetric keys depending on whether encryption is required and the number of auditors user  $i$  trusts.

**C<sub>i</sub>** is the digital signature over the fields of the same record  $i$  signed by the user  $U_i$  with key  $K_i^-$ , represented as:

$$C_i = \text{sign}_i(U_i, O_i, h(D_i), S_i)$$

Since the private key is confidential to a user, assuming it is not stolen, it is not possible to forge the signature of user  $i$ . The signature over the fields  $\langle U_i, O_i, h(D_i), S_i \rangle$  ensures the integrity of the record.

**P<sub>i</sub><sup>+</sup>** is the previous digital signature field which is a representation of the previous provenance record in the chain. The + indicates there may be more than one previous provenance record from different provenance chains. For the first user  $U_1$  in the provenance chain, this field is signed by the auditor with key  $K_{Aud}^-$ .

**N<sub>i</sub><sup>\*</sup>** is the next digital signature field which is a representation of the subsequent provenance record in the chain. The \* indicates there may be zero or more subsequent provenance record for a split into different provenance chains.



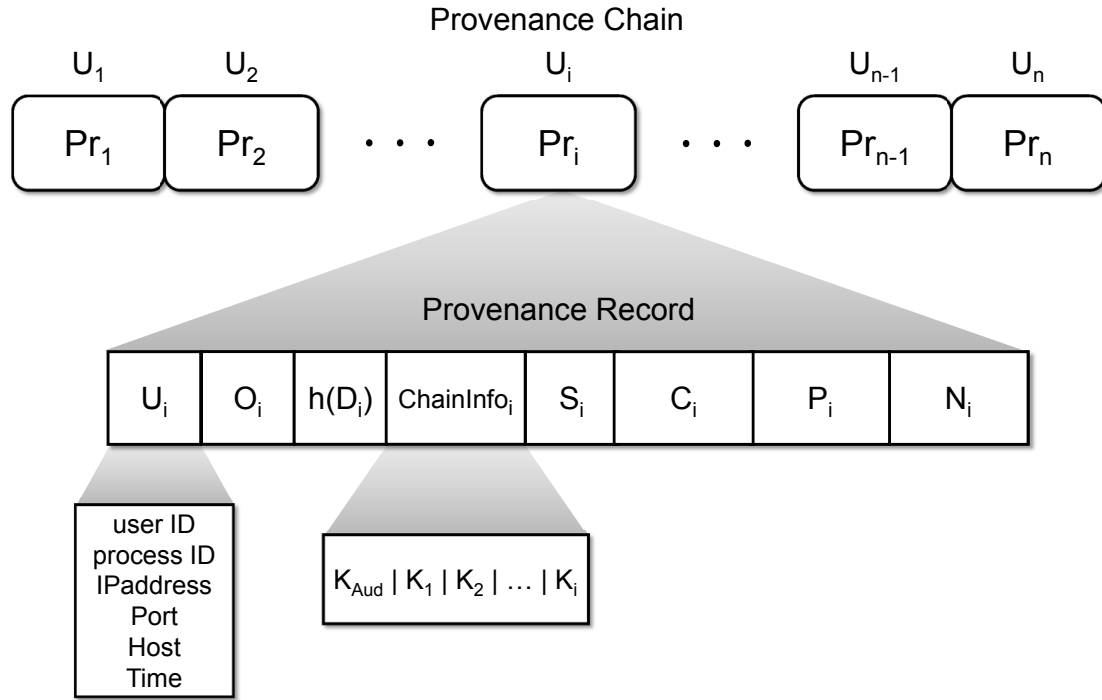


Figure 4.1. An illustration of the structure of provenance

The  $P^+$  and  $N^*$  fields are crucial for linking the different provenance records into a chain and for easing the verification process. They form the basis for the mutual agreement scheme. These fields are explained as:

For record  $i$ :

$$P_i = \text{sign}_{i-1}(h(D_{i-1}), ChainInfo_{i-1}|K_i, C_{i-1})$$

$$N_i = \text{sign}_{i+1}(h(D_i), ChainInfo_i|K_{i+1}, C_i)$$

For record  $i+1$ :

$$P_{i+1} = \text{sign}_i(h(D_i), ChainInfo_i|K_{i+1}, C_i)$$

$$N_{i+1} = \text{sign}_{i+2}(h(D_{i+1}), ChainInfo_{i+1}|K_{i+2}, C_{i+1})$$

It can be seen that the  $N$  field of record  $i$  is signed by user  $U_{i+1}$ . The  $P$  field of record  $i+1$  is signed by user  $U_i$ . User  $U_i$  after creating provenance record  $Pr_i$  passes the document to the user  $U_{i+1}$  who can then create the record  $Pr_{i+1}$ . An

---

**Algorithm 1** Provenance record creation steps
 

---

- 1: User  $i$  creates record  $Pr_i$ :
  - 2:  $Pr_i = \langle U_i, O_i, h(D_i), ChainInfo_i, S_i^*, C_i, P_i^+, N_i^* \rangle$
  - 3: **if**  $i = 1$  **then** ▷ User 1 is the creator of the document
  - 4:    $U_1 \leftarrow e_{S_1}(\text{User 1 Information})$
  - 5:    $O_1 \leftarrow \phi$  ▷  $U_1$  is the creator of the document
  - 6:    $h(D_1) \leftarrow \text{Hash of document created } D_1$
  - 7:    $ChainInfo_1 \leftarrow K_{Aud}|K_1$
  - 8:    $S_1 \leftarrow e_{Aud}(K_{S_1})$  ▷ Multiple for different auditors
  - 9:    $P_1 \leftarrow sign_{Aud}(IV, K_{Aud}|K_1, C_1)$  ▷ Auditor creates unique  $IV$  for this document
  - 10:    $N_1 \leftarrow \phi$
  - 11:    $C_1 \leftarrow sign_1(U_1, O_1, h(D_1), S_1)$
  - 12: **else** ▷ User  $i$  gets the document from user  $i-1$
  - 13:    $ChainInfo_i \leftarrow ChainInfo_{i-1}|K_i$
  - 14:    $P_i \leftarrow sign_{i-1}(h(D_{i-1}), ChainInfo_i, C_{i-1})$
  - 15:    $N_{i-1} \leftarrow sign_i(h(D_{i-1}), ChainInfo_i, C_{i-1})$
  - 16:   User  $i$  modifies document  $D_{i-1}$  to  $D_i$
  - 17:    $U_i \leftarrow e_{S_i}(\text{User } i \text{ Information})$
  - 18:    $O_i \leftarrow e_{S_i}(\text{Operations performed})$
  - 19:    $h(D_i) \leftarrow \text{Hash of modified document } D_i$
  - 20:    $S_i \leftarrow e_{Aud}(K_{S_i})$  ▷ Multiple for different auditors
  - 21:    $N_i \leftarrow \phi$
  - 22:    $C_i \leftarrow sign_i(U_i, O_i, h(D_i), S_i)$
  - 23: **end if**
-

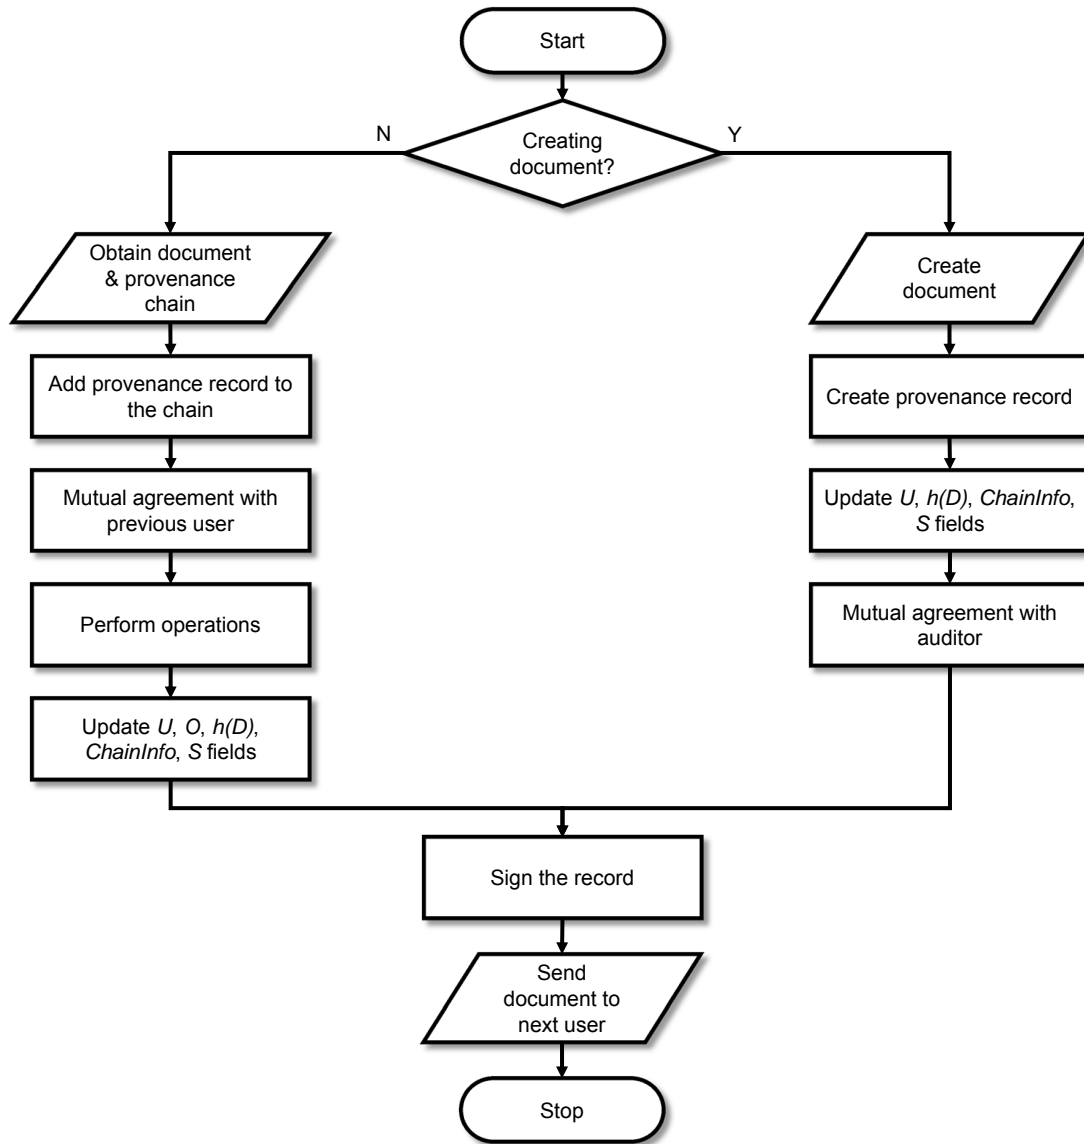


Figure 4.2. Flowchart of a user's actions for creating a provenance record

agreement is signed by both users, such that the record  $Pr_i$  contains the signature of  $U_{i+1}$  and  $Pr_{i+1}$  contains the signature of  $U_i$ . As can be seen, the fields  $N_i$  and  $P_{i+1}$  hold signatures over the same data which is the agreement between the users. The agreement between users  $U_i$  and  $U_{i+1}$  consists of the fields:

$$\langle h(D_i), ChainInfo_i | K_{i+1}, C_i \rangle$$

It indicates that the user  $U_i$  passes a document having hash  $h(D_i)$  to  $U_{i+1}$ , the sequence of users in the history of the document including  $U_{i+1}$  is  $ChainInfo|K_{i+1}$ , and  $C_i$  is the representation of the actions performed by  $h(D_i)$ . The agreement is an interactive hand-off mechanism where user  $U_i$  passes the document to user  $U_{i+1}$  and claims to have passed the provenance intact. It delineates the transition of responsibility of the document from user  $U_i$  to  $U_{i+1}$ .

The previous field for the first user in the provenance chain is signed by the auditor with key  $K_{Aud}$  and contains an initialization vector  $IV$  in the hash field which is known to the auditor. This  $IV$  is unique for each provenance chain. The purpose of  $IV$  and the  $P_1$  field signed by the auditor is to prevent an owner forgery attack which is discussed in detail in Chapter 5.  $IV$  is a place filler for the hash field, but is not required.

The steps followed by a user for creating a provenance record are described in Algorithm 1. It can be seen that the provenance record is constructed incrementally, the fields are created at each step when the user obtains the document, performs operations and passes it to the next user. Figure 4.2 shows a flowchart of the steps followed by a user for creating a provenance record. The user follows different courses of actions depending on whether he/she is creating the document, or is obtaining the document from another user. As can be seen, in both cases a mutual agreement takes place either with the auditor (if the user creates the document) or with the previous user (when the user receives the document).

### 4.3 Properties satisfied by our scheme

We briefly analyze our scheme for the properties discussed in Section 2.2.

**Confidentiality:** The information contained in the provenance records may be required to be kept confidential, for e.g., proprietary algorithms, identity of the user, etc. which are stored in the  $U$  and  $O$  fields. These fields need to be kept accessible only to the trusted auditor or group of trusted auditors. One approach is to





















an auditor can decrypt the symmetric keys stored in the record to further decrypt the  $U$  and  $O$  fields.

Figure 4.4 shows a flowchart of the steps followed by the auditor for auditing the provenance chain. The steps are similar to our discussion here. However, for the sake of simplicity of the flowchart, when any verification fails, the auditor enters the state 'Verify previous records' in which the verification of the remaining records up along the chain is performed to identify the cause of the verification failure.

In our scheme, each record is related only to its previous and next records. Unlike the Onion scheme, where every record contains the checksum information of the complete preceding chain, our scheme involves agreements with only pairs of records. Although Algorithm 2 involves checking the nodes sequentially from the last to the first, the verification can be conducted concurrently for pairs of records along the chain similar to the PKLC scheme. If all records are intact, this improves the verification process. If the verification fails for a particular pair of records, the preceding chain can then be investigated following the process described.

#### 4.6 Advantages of our scheme

We can see that the provenance records in our scheme are linked together into a chain (actually DAG) structure through the previous and next signature fields which serve as the mutual agreements between records. The scheme is simple to implement and has many advantages over the schemes discussed in Chapter 3. A summary of these advantages is given in Table 4.1.

First, it provides better protection than the Onion scheme against selective removal of provenance records from the chain. In the Onion scheme, records are loosely linked only in the forward direction, and a record does not contain any information about who the next or previous user(s) is. So, if an adversary receives a chain of records  $Pr_1|Pr_2|\dots|Pr_n$ , he/she can selectively remove part of the chain  $Pr_i|\dots|Pr_n$ , append a new record  $Pr_i$  with the signature over the signature of record  $Pr_{i-1}$  and















































integrity of individual records and the complete chain, which are common to all forms of provenance. We discuss some of the application scenarios here.

The scheme can be applied to employee review and recommendation systems. A document is passed between different employees, managers, teachers, etc. at different security levels and hierarchy. The document and provenance may have different sensitivity depending on the application and our scheme can support these scenarios. We club all such application scenarios as academic information sharing. Another scenario is a distributed application running instances on multiple nodes such that they incrementally process the data moving towards a sink. This is similar to a wireless sensor network, but encompasses a wider range of applications that can also include industrial pipelined processes. Our scheme can be applied for provenance generation in work-flows of industrial processes and grid based computing applications. In such scenarios, users maybe computer processes, sensor nodes or even industrial mechanized processes, which can have different representations in the provenance records. Wireless sensor networks are typically resource constrained and since our scheme employs computationally expensive operations, it may not be best suited for wireless sensor networks.

## 7.2 Future Work

Typical provenance schemes similar to our work in this thesis depend on transitive trust among users. A challenge that they face is the detection of collusion of users that are successive in the chain. In our future work we will look into a solution for detecting collusion of successive users. Another aspect is that our scheme requires a trusted auditor to prevent owner forgery cases and to access confidential information of provenance records. One direction of research is to reduce/remove the involvement of the trusted auditor and develop the scheme such that the users are the only entities involved. One possible way is to use Hierarchical Access Control (HAC) schemes such as [55] for efficient key management. In such a scheme, the creator of the document

is at the top of the hierarchy. This can remove the Public Key Infrastructure (PKI) required by most existing provenance schemes. However, it also introduces challenges of maintaining the desired security properties for the records and the chain. We will also work towards reducing the overhead of our scheme even further, by either employing faster algorithms, or by changing the procedure for mutual agreements. In addition, we will implement our scheme in a real world information sharing application, similar to ones discussed in the previous section and carry out a more thorough security and performance analysis.

## LIST OF REFERENCES

## LIST OF REFERENCES

- [1] Yogesh L. Simmhan, Beth Plale, and Dennis Gannon. A survey of data provenance in e-science. *SIGMOD Record*, 34(3):31–36, September 2005.
- [2] Luiz M. R. Gadelha Jr and Marta Mattoso. Kairos: An architecture for securing authorship and temporal information of provenance data in grid-enabled workflow management systems. In *Proceedings of the 2008 Fourth IEEE International Conference on eScience, ESCIENCE '08*, pages 597–602, Washington, DC, USA, 2008. IEEE Computer Society.
- [3] R. Hasan. Protecting the Past and Present of Data, with Applications in Provenance and Regulatory-Compliant Databases. *Proceedings of the Third SIGMOD PhD Workshop on Innovative Database Research (IDAR 2009)*, 2009.
- [4] James Cheney, Stephen Chong, Nate Foster, Margo Seltzer, and Stijn Vansumeren. Provenance: A future history. In *Proceedings of the 24th ACM SIGPLAN Conference Companion on Object Oriented Programming Systems Languages and Applications, OOPSLA '09*, pages 957–964, New York, NY, USA, 2009. ACM.
- [5] Shouhuai Xu, Qun Ni, Elisa Bertino, and Ravi Sandhu. A characterization of the problem of secure provenance management. In *Proceedings of the 2009 IEEE International Conference on Intelligence and Security Informatics, ISI'09*, pages 310–314, Piscataway, NJ, USA, 2009. IEEE Press.
- [6] Wang chiew Tan. Research problems in data provenance. *IEEE Data Engineering Bulletin*, 27:45–52, 2004.
- [7] Rajendra Bose and James Frew. Lineage retrieval for scientific data processing: A survey. *ACM Computing Surveys (CSUR)*, 37(1):1–28, March 2005.
- [8] Simon Miles, Sylvia C. Wong, Weijian Fang, Paul Groth, Klaus peter Zauner, and Luc Moreau. Provenance-based validation of e-science experiments. In *In ISWC*, pages 801–815. Springer-Verlag, 2005.
- [9] Peter Buneman, Adriane Chapman, and James Cheney. Provenance management in curated databases. In *Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data, SIGMOD '06*, pages 539–550, New York, NY, USA, 2006. ACM.
- [10] James Cheney, Laura Chiticariu, and Wang-Chiew Tan. Provenance in databases: Why, how, and where. *Found. Trends databases*, 1(4):379–474, April 2009.

- [11] Paolo Missier, Saumen Dey, Khalid Belhajjame, Víctor Cuevas-Vicenttín, and Bertram Ludäscher. D-prov: Extending the prov provenance model with workflow structure. In *Proceedings of the 5th USENIX Workshop on the Theory and Practice of Provenance*, TaPP '13, pages 9:1–9:7, Berkeley, CA, USA, 2013. USENIX Association.
- [12] Fernando Chirigati, Dennis Shasha, and Juliana Freire. Reprozip: Using provenance to support computational reproducibility. In *Proceedings of the 5th USENIX Workshop on the Theory and Practice of Provenance*, TaPP '13, pages 1:1–1:4, Berkeley, CA, USA, 2013. USENIX Association.
- [13] S. Jensen, B. Plale, M.S. Aktas, Yuan Luo, Peng Chen, and H. Conover. Provenance capture and use in a satellite data processing pipeline. *Geoscience and Remote Sensing, IEEE Transactions on*, 51(11):5090–5097, Nov 2013.
- [14] M.R. Huq, P.M.G. Apers, and A. Wombacher. An inference-based framework to manage data provenance in geoscience applications. *Geoscience and Remote Sensing, IEEE Transactions on*, 51(11):5113–5130, Nov 2013.
- [15] Nicholas E Taylor and Zachary G Ives. Reconciling while tolerating disagreement in collaborative data sharing. In *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*, pages 13–24. ACM, 2006.
- [16] Todd J. Green, Grigoris Karvounarakis, Zachary G. Ives, and Val Tannen. Update exchange with mappings and provenance. In *Proceedings of the 33rd International Conference on Very Large Data Bases*, VLDB '07, pages 675–686. VLDB Endowment, 2007.
- [17] Todd J. Green, Grigoris Karvounarakis, Nicholas E. Taylor, Olivier Biton, Zachary G. Ives, and Val Tannen. Orchestra: Facilitating collaborative data sharing. In *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data*, SIGMOD '07, pages 1131–1133, New York, NY, USA, 2007. ACM.
- [18] Zachary G. Ives, Todd J. Green, Grigoris Karvounarakis, Nicholas E. Taylor, Val Tannen, Partha Pratim Talukdar, Marie Jacob, and Fernando Pereira. The orchestra collaborative data sharing system. *SIGMOD Rec.*, 37(3):26–32, September 2008.
- [19] John Ockerbloom. Copyright and provenance: Some practical problems. *IEEE Data Eng. Bull.*, 30(4):51–58, 2007.
- [20] Ragib Hasan, Radu Sion, and Marianne Winslett. Introducing secure provenance: Problems and challenges. In *Proceedings of the 2007 ACM Workshop on Storage Security and Survivability*, StorageSS '07, pages 13–18, New York, NY, USA, 2007. ACM.
- [21] Kiran-Kumar Muniswamy-Reddy, Peter Macko, and Margo Seltzer. Making a cloud provenance-aware. In *First Workshop on on Theory and Practice of Provenance*, TAPP'09, pages 12:1–12:10, Berkeley, CA, USA, 2009. USENIX Association.
- [22] Kiran-Kumar Muniswamy-Reddy and Margo Seltzer. Provenance as first class cloud data. *SIGOPS Oper. Syst. Rev.*, 43(4):11–16, January 2010.

- [23] Kiran-Kumar Muniswamy-Reddy, Peter Macko, and Margo Seltzer. Provenance for the cloud. In *Proceedings of the 8th USENIX Conference on File and Storage Technologies*, FAST'10, pages 15–14, Berkeley, CA, USA, 2010. USENIX Association.
- [24] Boris Glavic, Javed Siddique, Periklis Andritsos, and Renée J. Miller. Provenance for data mining. In *Proceedings of the 5th USENIX Workshop on the Theory and Practice of Provenance*, TaPP '13, pages 5:1–5:5, Berkeley, CA, USA, 2013. USENIX Association.
- [25] Nathaniel Husted, Sharjeel Qureshi, Dawood Tariq, and Ashish Gehani. Android provenance: Diagnosing device disorders. In *Proceedings of the 5th USENIX Workshop on the Theory and Practice of Provenance*, TaPP '13, pages 4:1–4:7, Berkeley, CA, USA, 2013. USENIX Association.
- [26] Boris Glavic, Kyumars Sheykh Esmaili, Peter Michael Fischer, and Nesime Tatbul. Ariadne: Managing fine-grained provenance on data streams. In *Proceedings of the 7th ACM International Conference on Distributed Event-based Systems*, DEBS '13, pages 39–50, New York, NY, USA, 2013. ACM.
- [27] S. Sultana, M. Shehab, and E. Bertino. Secure provenance transmission for streaming data. *Knowledge and Data Engineering, IEEE Transactions on*, 25(8):1890–1903, Aug 2013.
- [28] Uri Braun, Avraham Shinnar, and Margo Seltzer. Securing provenance. In *Proceedings of the 3rd Conference on Hot Topics in Security*, HOTSEC'08, pages 4:1–4:5, Berkeley, CA, USA, 2008. USENIX Association.
- [29] Ragib Hasan, Radu Sion, and Marianne Winslett. Preventing history forgery with secure provenance. *ACM Transactions on Storage (TOS)*, 5(4):12:1–12:43, December 2009.
- [30] Xinlei Wang, Kai Zeng, K. Govindan, and P. Mohapatra. Chaining for securing data provenance in distributed information networks. In *Military Communications Conference, 2012 - MILCOM 2012*, pages 1–6, Orlando, FL, USA, Oct 2012.
- [31] Paul Groth, Michael Luck, and Luc Moreau. A protocol for recording provenance in service-oriented grids. In *In Proceedings of the 8th International Conference on Principles of Distributed Systems (OPODIS04)*, pages 124–139, 2004.
- [32] Peter Buneman, Sanjeev Khanna, and Wang Chiew Tan. Why and where: A characterization of data provenance. In *Proceedings of the 8th International Conference on Database Theory*, ICDT '01, pages 316–330, London, UK, UK, 2001. Springer-Verlag.
- [33] Kiran-Kumar Muniswamy-Reddy, David A Holland, Uri Braun, and Margo I Seltzer. Provenance-aware storage systems. In *USENIX Annual Technical Conference, General Track*, pages 43–56, 2006.
- [34] Adriane P. Chapman, H. V. Jagadish, and Prakash Ramanan. Efficient provenance storage. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, SIGMOD '08, pages 993–1006, New York, NY, USA, 2008. ACM.



- [35] Ryan K. L. Ko, Peter Jagadpramana, and Bu Sung Lee. Flogger: A file-centric logger for monitoring file access and transfers within cloud computing environments. In *Proceedings of the 2011 IEEE 10th International Conference on Trust, Security and Privacy in Computing and Communications, TRUSTCOM '11*, pages 765–771, Washington, DC, USA, 2011. IEEE Computer Society.
- [36] Salmin Sultana and Elisa Bertino. A file provenance system. In *Proceedings of the Third ACM Conference on Data and Application Security and Privacy, CODASPY '13*, pages 153–156, New York, NY, USA, 2013. ACM.
- [37] Lucian Carata, Ripduman Sohan, Andrew Rice, and Andy Hopper. Ipapi: Designing an improved provenance api. In *Proceedings of the 5th USENIX Workshop on the Theory and Practice of Provenance, TaPP '13*, pages 10:1–10:4, Berkeley, CA, USA, 2013. USENIX Association.
- [38] Devdatta Kulkarni. A provenance model for key-value systems. In *Proceedings of the 5th USENIX Workshop on the Theory and Practice of Provenance, TaPP '13*, pages 12:1–12:4, Berkeley, CA, USA, 2013. USENIX Association.
- [39] Nikilesh Balakrishnan, Thomas Bytheway, Ripduman Sohan, and Andy Hopper. Opus: A lightweight system for observational provenance in user space. In *Proceedings of the 5th USENIX Workshop on the Theory and Practice of Provenance, TaPP '13*, pages 8:1–8:4, Berkeley, CA, USA, 2013. USENIX Association.
- [40] Sherif Akoush, Ripduman Sohan, and Andy Hopper. Hadoopprov: Towards provenance as a first class citizen in mapreduce. In *Proceedings of the 5th USENIX Workshop on the Theory and Practice of Provenance, TaPP '13*, pages 11:1–11:4, Berkeley, CA, USA, 2013. USENIX Association.
- [41] Yulai Xie, Kiran-Kumar Muniswamy-Reddy, Dan Feng, Yan Li, and Darrell D. E. Long. Evaluation of a hybrid approach for efficient provenance storage. *Trans. Storage*, 9(4):14:1–14:29, November 2013.
- [42] Zachary Hensley, Jibonananda Sanyal, and Joshua New. Provenance in sensor data management. *Commun. ACM*, 57(2):55–62, February 2014.
- [43] Victor Tan, Paul Groth, Simon Miles, Sheng Jiang, Steve Munroe, and Luc Moreau. Security issues in a soa-based provenance system. In *In Proceedings of the International Provenance and Annotation Workshop (IPAW06)*. Springer-Verlag, 2006.
- [44] Muhammad Imran and Helmut Hlavacs. Provenance in the cloud: Why and how? In *CLOUD COMPUTING 2012, The Third International Conference on Cloud Computing, GRIDs, and Virtualization*, pages 106–112, 2012.
- [45] Rongxing Lu, Xiaodong Lin, Xiaohui Liang, and Xuemin (Sherman) Shen. Secure provenance: The essential of bread and butter of data forensics in cloud computing. In *Proceedings of the 5th ACM Symposium on Information, Computer and Communications Security, ASIACCS '10*, pages 282–292, New York, NY, USA, 2010. ACM.
- [46] Ryan K. L. Ko, Peter Jagadpramana, Miranda Mowbray, Siani Pearson, Markus Kirchberg, Qianhui Liang, and Bu Sung Lee. Trustcloud: A framework for accountability and trust in cloud computing. In *Proceedings of the 2011 IEEE World Congress on Services, SERVICES '11*, pages 584–588, Washington, DC, USA, 2011. IEEE Computer Society.

- [47] Adam Bates, Ben Mood, Masoud Valafar, and Kevin Butler. Towards secure provenance-based access control in cloud environments. In *Proceedings of the Third ACM Conference on Data and Application Security and Privacy, CODASPY '13*, pages 277–284, New York, NY, USA, 2013. ACM.
- [48] Uri Braun and Avi Shinnar. A security model for provenance. Technical report, Computer Science Group, Harvard University, 2006.
- [49] S. Sultana, G. Ghinita, E. Bertino, and M. Shehab. A lightweight secure provenance scheme for wireless sensor networks. In *Parallel and Distributed Systems (ICPADS), 2012 IEEE 18th International Conference on*, pages 101–108, Dec 2012.
- [50] S. Sultana, G. Ghinita, E. Bertino, and M. Shehab. A lightweight secure scheme for detecting provenance forgery and packet drop attacks in wireless sensor networks. *Dependable and Secure Computing, IEEE Transactions on*, PP(99):1–1, 2013.
- [51] K. Alharbi and Xiaodong Lin. Pdp: A privacy-preserving data provenance scheme. In *Distributed Computing Systems Workshops (ICDCSW), 2012 32nd International Conference on*, pages 500–505, June 2012.
- [52] Hasan, Ragib and Sion, Radu and Winslett, Marianne. The case of the fake picasso: Preventing history forgery with secure provenance. In *Proceedings of the 7th Conference on File and Storage Technologies, FAST '09*, pages 1–14, Berkeley, CA, USA, 2009. USENIX Association.
- [53] Amril Syalim, Takashi Nishide, and Kouichi Sakurai. Preserving integrity and confidentiality of a directed acyclic graph model of provenance. In *Proceedings of the 24th Annual IFIP WG 11.3 Working Conference on Data and Applications Security and Privacy, DBSec'10*, pages 311–318, Berlin, Heidelberg, 2010. Springer-Verlag.
- [54] Deepak Ajwani, Adan Cosgaya-Lozano, and Norbert Zeh. A topological sorting algorithm for large graphs. *J. Exp. Algorithmics*, 17:3.2:3.1–3.2:3.21, September 2012.
- [55] Mikhail J. Atallah, Marina Blanton, Nelly Fazio, and Keith B. Frikken. Dynamic and efficient key management for access hierarchies. *ACM Trans. Inf. Syst. Secur.*, 12(3):18:1–18:43, January 2009.