Personal data utilization and protection through algorithms and protocols for privacy-preserving electronic transactions

Aimilia Tasidou

Department of Electrical and Computer Engineering

Democritus University of Thrace

Advisor: Pavlos S. Efraimidis

A thesis submitted for the degree of

Doctor of Philosophy

Xanthi, October 2015

Democritus University of Thrace
Department of Electrical and Computer Engineering
Building A, ECE, University Campus - Kimmeria, 67100 Xanthi, Greece

I would like to dedicate this thesis to my parents.

Acknowledgements

In the course of my PhD work, I have received the support of several people, who have contributed directly or indirectly to the completion of this thesis. Therefore, I would like acknowledge their contributions and express my gratitude to them.

First, I would like to express my appreciation and thanks to my advisor Prof. Pavlos Efraimidis, for giving me the opportunity of this PhD candidacy and for supporting me in the course of my research, contributing his guidance, time, ideas and practical advice.

I would also like to thank the rest of my advisory committee members, Prof. Alexandros S. Karakos and Prof. Sokratis Katsikas, for promptly attending to all necessary academic and administrative issues with respect to my PhD studies. Additionally, I would like to express my gratitude to the members of my examining committee, Prof. Vasilis Katos, Prof. Lilian Mitrou, Prof. Avi Arampatzis and Prof. Eli Katsiri, for the evaluation of my work.

Thanks also belong to the members of the SPHINX and myVisitPlanner research projects, within which some of the works included in this thesis were developed. I am particularly grateful to Prof. Vasilis Katos, Prof. Lilian Mitrou, Dr. George Drosatos, Dr. Yiannis Soupionis, Aris Koutsiamanis and Fotis Nalbadis for their the constructive collaboration in preparing joint works described in this dissertation. I would like to address special thanks to Prof. Vasilis Katos for the constructive and motivating discussions on my research. Additionally, I would like to express my gratitude to Dr. Christos Emmanouilidis for his support and encouragement during our collaboration, as well as for taking an interest in my research topic. A special thanks also goes to my research group colleague and good friend George Drosatos for his excellent collaboration, helpfulness and company along the way.

I would also like to take the opportunity to thank and acknowledge the influence of the people from my previous studies, who inspired and encouraged me to pursue academic research, Prof. Themis Panayiotopoulos, Prof. Panagiotis Tsikouras and Prof. Michael Rovatsos.

Furthermore, I would like to express my gratefulness to my partner in life, friend and colleague Aris Koutsiamanis, for travelling this path with me and for his patience, understanding and encouragement, as well as his practical help and constructive collaboration on academic matters.

Finally, a special thanks to my family, as well as Aris' family, for their patience, love, practical and moral support and constant encouragement.

Abstract

The widespread use of information and communication technologies and electronic transactions today leads to vast amounts of personal data being gathered and stored, making the need for personal data protection more imminent than ever. The protection of personal privacy and the negative externalities that arise from the exploitation of personal information have become growing concerns for Internet users.

Moreover, transactions today are conducted in a way that leaves the consumer at a disadvantage when it comes to privacy protection. Sensitive private information is left at the disposal of companies and is often (un)intentionally leaked to unauthorised parties. As a result, incidents of private data loss are an almost daily occurrence. There is a growing demand for privacy-preserving management of private information that will make individuals feel safer during their transactions and assist companies with customer data management.

Additionally, personal data processing brings about important benefits to the marketplace for both companies and customers. Therefore, a fair, effective and legitimate way to acquire individuals' personal information for processing should exist. Indeed, the aspiration of privacy protection technology is not to lock all personal information away from any possible access or use, but to allow access to personal information in a controlled manner. Moreover, it is possible that individuals would be willing to provide access to (some of) their personal information in exchange for some profit, provided they could be assured of the safety of their information.

In this thesis we investigate and propose methods for the utilization and protection of personal information, offering data owners the ability to maintain the benefits of the use of their information in electronic transactions, while retaining control over the way their information is distributed and used. We argue that today the enabling technologies exist to create such accountable privacy-preserving services, protecting parties from fraud and identity theft. Following the principle that

personal data should remain at its owner's control, we propose privacy-preserving applications that allow data owners to protect and utilize their information to receive personalised services without giving up their privacy.

More specifically, starting from the Polis personal data management framework, we proceed to propose applications that address the economic aspects of privacy, the FPIT personal information market and the Portfolio consumer data management system. We apply the idea of FPIT in the PrivTAM system for privacy-preserving television audience measurements and the Portfolio concept in the PrivCAPTCHA system for user-centric VoIP CAPTCHA adaptation. Experimental implementations are developed in several of the proposed applications to illustrate the feasibility of the proposed solutions. Moreover, privacy protection principles are applied to achieve user profile data protection in myVisitPlanner, a real-world travel recommendation and planning system.

Extended Abstract in Greek (Περίληψη)

Ο τρόπος που διεξάγονται οι ηλεκτρονικές συναλλαγές σήμερα αναγκάζει τους πολίτες να παραχωρούν τα προσωπικά τους δεδομένα σε εταιρείες και φορείς, χωρίς να γνωρίζουν πώς αυτά θα χρησιμοποιηθούν στην συνέχεια. Περιστατικά σκόπιμης ή μη διαρροής προσωπικών δεδομένων έρχονται στο φως της δημοσιότητας καθημερινά, προκαλώντας την έντονη ανησυχία των πολιτών.

Η ψηφιοποίηση πολλών δραστηριοτήτων της καθημερινότητας του σύγχρονου ατόμου, οι αυξημένες δυνατότητες αποθήκευσης μεγάλου όγκου δεδομένων και τα ισχυρά μέσα επεξεργασίας που υπάρχουν διαθέσιμα σήμερα, καθιστούν τη διαρροή προσωπικών δεδομένων σημαντική απειλή για την ασφάλεια και την ευημερία των πολιτών. Η έλλειψη προστασίας της ιδιωτικότητας και οι αρνητικές συνέπειες που προκύπτουν από την καταχρηστική εκμετάλλευση των προσωπικών δεδομένων, απασχολεί όλο και περισσότερο τους σύγχρονους πολίτες, οι οποίοι νιώθουν ανίκανοι να ελέγξουν πού διακινούνται και με ποιόν τρόπο χρησιμοποιούνται τα προσωπικά τους δεδομένα. Το γεγονός αυτό έχει άμεσες επιπτώσεις στην εμπιστοσύνη των πολιτών προς τις τεχνολογίες της Πληροφορικής και αποτελεί εμπόδιο στην ανάπτυξη και διάδοση εφαρμογών ηλεκτρονικών δραστηριοτήτων.

Από την άλλη πλευρά, οι δυνατότητες μαζικής συλλογής προσωπικών δεδομένων από τις όλο και αυξανόμενες ηλεκτρονικές δραστηριότητες των πολιτών, ανοίγει νέες ευκαιρίες αξιοποίησης των δεδομένων αυτών, τόσο προς όφελος των εταιρειών, όσο και προς όφελος του ίδιου του πολίτη, ο οποίος έχει πλέον την δυνατότητα να μελετήσει την ίδια του τη συμπεριφορά και τις συνήθειές του και να εξάγει χρήσιμα συμπεράσματα από τα δεδομένα αυτά. Επιπλέον, η δυνατότητα χρήσης των δεδομένων των χρηστών για την αποτελεσματικότερη εξατομίκευση των υπηρεσιών που τους παρέχονται αποτελεί σημαντικό παράγοντα βελτίωσης των παρεχόμενων υπηρεσιών σήμερα.

Για τους παραπάνω λόγους, καθίσταται κρίσιμη η ανάπτυξη σύγχρονων συστημάτων προστασίας της ιδιωτικότητας, τα οποία θα ενισχύουν την σωστότερη και αποδοτικότερη αξιοποίηση των προσωπικών δεδομένων, μειώνοντας τις πιθανότητες διαρροής και των κινδύνων που αυτή ενέχει.

Η οικονομική διάσταση της ιδιωτικότητας (Economics of Privacy) αποτελεί ένα σύγχρονο και ιδιαίτερα επίκαιρο αντικείμενο έρευνας. Η επεξεργασία προσωπικών δεδομένων επιφέρει σημαντικά κέρδη στην αγορά σήμερα, αλλά ταυτόχρονα κρύβει σημαντικούς κινδύνους για την ιδιωτικότητα των πολιτών, με σημαντικά κόστη, σχετικά τόσο με την ασφάλεια όσο και με την οικονομία. Ταυτόχρονα, ο πολίτης δεν λαμβάνει κανένα μέρος από τα κέρδη των εταιρειών από την εκμετάλλευση των προσωπικών του δεδομένων. Για τους παραπάνω λόγους, ενισχύεται η ιδέα ότι ο πολίτης έχει το δικαίωμα να γνωρίζει και να ελέγχει τους αποδέκτες και τις χρήσεις των προσωπικών του δεδομένων, καθώς και να αποζημιώνεται για την χρήση τους.

Η παρούσα διδακτορική έρευνα περιλαμβάνει τη διερεύνηση των τρόπων αξιοποίησης και προστασίας προσωπικών δεδομένων προς όφελος των κατόχων τους, προσφέροντάς τους τον έλεγχο της διακίνησης των προσωπικών τους πληροφοριών, καθώς και οφέλη από την χρήση τους στις συναλλαγές τους. Βασική αρχή της έρευνας είναι η ιδέα ότι τα προσωπικά δεδομένα των πολιτών οφείλουν να παραμένουν στον έλεγχο των ίδιων των κατόχων τους, με συνέπεια αυτοί να παραμένουν ενήμεροι για τους σκοπούς, αλλά και για τους τρόπους που αυτά χρησιμοποιούνται. Στόχος είναι η δημιουργία ενός δίκαιου, ασφαλούς και νόμιμου τρόπου απόκτησης πρόσβασης και αξιοποίησης των προσωπικών δεδομένων των πολιτών. Άλλωστε, ο σκοπός της προστασίας της ιδιωτικότητας δεν είναι να αποκλειστεί η πρόσβαση σε κάθε είδους προσωπικά δεδομένα, αλλά να επιτευχθεί η πρόσβαση στα κατάλληλα δεδομένα με ελεγχόμενο τρόπο.

Με βάση τις παραπάνω επιδιώξεις, προτάθηκαν και υλοποιήθηκαν οι παρακάτω εργασίες που συντελούν στην αξιοποίηση των δεδομένων των χρηστών, διασφαλίζοντας ταυτόχρονα την ιδιωτικότητά τους:

- Η πλατφόρμα διαχείρισης προσωπικών δεδομένων Polis για χρήση κατά τη διεξαγωγή ηλεκτρονικών συναλλαγών, με σκοπό την διαφύλαξη της ιδιωτικότητας των προσωπικών δεδομένων του χρήστη.
- Η «αγορά» προσωπικών δεδομένων (market for personal information) FPIT (Fair Personal Information Trades), όπου παρέχεται ελεγχόμενη πρόσβαση στα προσωπικά δεδομένα ιδιωτών για ενδιαφερόμενους φορείς και εταιρείες, προσφέροντας ταυτόχρονα δίκαιη αποζημίωση στους κατόχους των δεδομένων μέσω τεχνικών μικροπληρωμών (micropayments).
- Το σύστημα μετρήσεων τηλεθέασης PrivTAM με προστασία της ιδιωτικότητας των συμμετεχόντων, με εφαρμογή τεχνικών δίκαιων συναλλαγών προσωπικών πληροφοριών (Fair Personal Information Trades)

και ασφαλών υπολογισμών.

- Ο ηλεκτρονικός χαρτοφύλακας (Portfolio) για την ελεγχόμενη αποθήκευση, διαχείριση και αξιοποίηση των προσωπικών δεδομένων του χρήστη σε συναλλαγές (ηλεκτρονικές και φυσικές). Ανάπτυξη πρωτοκόλλων για την προστασία των προσωπικών δεδομένων του χρήστη (unlinkability, privacy) κατά την παροχή ελεγχόμενης πρόσβασης σε αυτά και την ταυτόχρονη εξασφάλιση της ακεραιότητας και της εγκυρότητας των δεδομένων (accountability).
- Το σύστημα PrivCAPTCHA προσαρμογής των προκλήσεων ηχητικών CAPTCHA σύμφωνα με τα χαρακτηριστικά και τις ιδιότητες του χρήστη, με ταυτόχρονη παροχή προστασίας της ιδιωτικότητάς του, που κάνει χρήση των τεχνολογιών και των μεθόδων που αναπτύχθηκαν στο Portfolio.
- Το υποσύστημα προστασίας των δεδομένων των προφίλ χρηστών στο σύστημα συστάσεων και χρονοπρογραμματισμού ταξιδιών myVisitPlanner.

Οι εργασίες αυτές αποδεικνύουν ότι είναι εφικτή η αξιοποίηση και ταυτόχρονα η προστασία των προσωπικών δεδομένων των χρηστών. Συνοπτικές περιγραφές των εργασιών δίνονται παρακάτω στις περιλήψεις των αντίστοιχων κεφαλαίων της διατριβής.

ΔΟΜΗ ΤΗΣ ΔΙΑΤΡΙΒΗΣ

Παρακάτω συνοψίζονται τα περιεχόμενα των κεφαλαίων της διατριβής:

Κεφάλαιο 1: Introduction

Περιγράφονται τα κίνητρα και οι στόχοι της παρούσας διατριβής. Περιγράφεται η σημερινή κατάσταση σε σχέση με τις ηλεκτρονικές συναλλαγές και την προστασία των προσωπικών δεδομένων κατά τις συναλλαγές (ηλεκτρονικές και μη), καθώς και οι δυνατότητες αξιοποίησης των προσωπικών δεδομένων για την βελτίωση των παρεχόμενων υπηρεσιών προς τους κατόχους τους.

Κεφάλαιο 2: Background

Παρέχεται το απαραίτητο υπόβαθρο για τη κατανόηση των βασικών εννοιών που χρησιμοποιούνται σε αυτή την διατριβή. Πιο συγκεκριμένα, περιγράφονται έννοιες που αφορούν την ιδιωτικότητα, την ασφάλεια και

τη κρυπτογραφία και χρησιμοποιούνται ως δομικά στοιχεία στην ανάπτυξη των προτεινόμενων λύσεων της διατριβής.

Κεφάλαιο 3: The Polis Personal Data Management Framework

Περιγράφεται η πλατφόρμα διαχείρισης προσωπικών δεδομένων Polis, που σκοπό έχει την ενίσχυση της προστασίας της ιδιωτικότητας κατά τις ηλεκτρονικές συναλλαγές, παραχωρώντας στον χρήστη τη διαχείριση των προσωπικών του δεδομένων. Η προσέγγιση αυτή μπορεί να επιφέρει οφέλη τόσο στους χρήστες, όσο και στις επιχειρήσεις με τις οποίες συναλλάσσονται.

Προτείνεται, αναπτύσσεται και αξιολογείται η αρχιτεκτονική διαχείρισης προσωπικών δεδομένων Polis. Περιγράφονται αντιπροσωπευτικές ηλεκτρονικές συναλλαγές που αφορούν δεδομένα προσωπικού χαρακτήρα και προτείνονται πρωτόκολλα διεξαγωγής των συναλλαγών αυτών στα πλαίσια του συστήματος Polis. Οι προτεινόμενες λύσεις δοκιμάζονται τόσο σε πρωτότυπο ως πιλοτική εφαρμογή, όσο και ως μέρος ενός εμπορικού συστήματος διαχείρισης δεδομένων. Τα αποτελέσματα αυτής της εργασίας δείχνουν ότι οι ηλεκτρονικές συναλλαγές μπορούν να παραμείνουν εφικτές και απλές, διατηρώντας τα προσωπικά δεδομένα μόνο στην πλευρά του ιδιοκτήτη τους.

Κεφάλαιο 4: Fair Personal Information Trades

Περιγράφεται το σύστημα FPIT (Fair Personal Information Trades), το οποίο διερευνά τη δυνατότητα δημιουργίας μίας «αγοράς» προσωπικών δεδομένων (market for personal information), η οποία θα επιτρέπει σε ενδιαφερόμενους να αποκτούν ελεγχόμενη πρόσβαση σε δεδομένα χρηστών, παρέχοντας αντίστοιχη αμοιβή για την πρόσβαση αυτή (οικονομική ή προνομίων). Τα προσωπικά δεδομένα βρίσκονται αποκλειστικά στην κατοχή των ιδιοκτητών τους και δεν επιτρέπεται η αποθήκευσή τους στην πλευρά των εταιρειών. Η πρόσβαση στα προσωπικά δεδομένα του πολίτη επιτυγχάνεται με τη χρήση κατάλληλων αδειών πρόσβασης (data licences), οι οποίες δίδονται από τον πολίτη έπειτα από σύναψη κατάλληλης συμφωνίας με τον ενδιαφερόμενο. Ο τρόπος και η διάρκεια χρήσης των δεδομένων ορίζεται αυστηρά κατά της σύναψη της συμφωνίας. Καθώς τα προσωπικά δεδομένα δεν επιτρέπεται να αποθηκεύονται από την εταιρεία, η άδεια πρόσβασης δίνει την δυνατότητα στην εταιρεία να ανακτά τα επιθυμητά δεδομένα κάθε φορά που χρειάζεται πρόσβαση σε αυτά. Με τον τρόπο αυτό, ο πολίτης μπορεί να διατηρεί τον έλεγχο τον δεδομένων

του και να παραμένει ενήμερος για κάθε ανάκτηση τους. Η υλοποίηση του συστήματος FPIT γίνεται μέσω ενσωμάτωσης του συστήματος διαχείρισης προσωπικών δεδομένων Polis. Επίσης ενσωματώνεται υποσύστημα μικροπληρωμών (micropayments) το οποίο επιτρέπει την αποτελεσματική και αποδοτική πραγματοποίηση μεγάλου όγκου συναλλαγών μικρής αξίας.

Κεφάλαιο 5: Privacy-preserving Television Audience Measurements

Το σύστημα PrivTAM αποτελεί μία νέα προσέγγιση για την επίτευξη μετρήσεων τηλεθέασης (Television Audience Measurement – TAM), η οποία αξιοποιώντας τις δυνατότητες των Smart TV τεχνολογιών, επιτρέπει την πραγματοποίηση εγκυρότερων μετρήσεων, σε μεγαλύτερο δείγμα τηλεθεατών, προστατεύοντας ταυτόχρονα την ιδιωτικότητα των δεδομένων των συμμετεχόντων. Η χρήση κρυπτογραφικών τεχνικών κατά τον υπολογισμό των μετρήσεων τηλεθέασης εξασφαλίζει τόσο η προστασία της ιδιωτικότητας των συμμετεχόντων όσο και την εγκυρότητα των αποτελεσμάτων. Επιπλέον, στους συμμετέχοντες δίνεται η δυνατότητα να αποζημιωθούν για την συμμετοχή τους στη μέτρηση και τα δεδομένα που παρείχαν, επιτρέποντας έτσι να επιστραφεί στους κατόχους των δεδομένων μέρος των κερδών που προκύπτει από τη δραστηριότητα των εταιρειών τηλεθέασης και παραγωγής τηλεοπτικών προγραμμάτων. Τα πειραματικά αποτελέσματα σε Android-based Smart TVs επιβεβαιώνουν τη βιωσιμότητα της προσέγγισης.

Κεφάλαιο 6: A Personal Data Management Portfolio for Privacy-preserving Consumer Data Utilization

Σε αυτό το κεφάλαιο ορίζεται το σύστημα χαρτοφύλακα (Portfolio) προσωπικών δεδομένων, μέσω του οποίου οι χρήστες μπορούν να προστατεύουν, αλλά και ταυτόχρονα να αξιοποιούν τα προσωπικά τους δεδομένα με σκοπό την καλύτερη εξυπηρέτησή τους, αλλά και την απόκτηση εξατομικευμένων προσφορών σε σχέση με το προφίλ και τα χαρακτηριστικά τους.

Προβλέπεται η χρήση του συστήματος για κάθε είδους προσωπικά δεδομένα, συμπεριλαμβανομένων δεδομένων προτίμησης και συμπεριφοράς, όπως και αγοραστικού ιστορικού. Για παράδειγμα, η εφαρμογή αυτή θα μπορούσε να αντικαταστήσει την πολιτική διατήρησης προφίλ χρηστών με τις αγορές και τις προτιμήσεις τους σε εμπορικά καταστήματα και super markets, καθιστώντας δυνατή την επιλογή του χρήστη να συμμετέχει ή όχι και σε ποιό βαθμό σε έρευνες που αφορούν προσωπικά του

δεδομένα. Επίσης μία τέτοια εφαρμογή θα μπορούσε να εξυπηρετήσει την παροχή τέτοιου είδους πληροφορίας σε πολλαπλούς αποδέκτες, για παράδειγμα οι αγορές του χρήστη στην μία αλυσίδα καταστημάτων να παρέχονται προς επεξεργασία σε μία άλλη, πάντα με την συγκατάθεση του χρήστη και με δίκαιη αποζημίωσή του. Ο χαρτοφύλακας του κάθε χρήστη βρίσκεται υπό το δική του διαχείριση και αποθηκεύεται είτε τοπικά σε υπολογιστική υποδομή του χρήστη είτε κρυπτογραφημένο σε cloud υπηρεσίες μέσω Internet.

Εφόσον τα δεδομένα βρίσκονται στον έλεγχο του πολίτη, προκύπτουν θέματα εγκυρότητας των δεδομένων. Προτείνονται μηχανισμοί διασφάλισης της εγκυρότητας και της αξιοπιστίας των δεδομένων του χαρτοφύλακα. Αναπτύσσονται αλγόριθμοι και πρωτόκολλα για την εξασφάλιση της ακεραιότητας και της εγκυρότητας των δεδομένων που υποβάλλονται από τους χρήστες.

Κεφάλαιο 7: Privacy-preserving, User-centric VoIP CAPTCHA Challenges

Περιγράφεται το σύστημα PrivCAPTCHA προσαρμογής των προκλήσεων ηχητικών CAPTCHA σύμφωνα με τα χαρακτηριστικά και τις ιδιότητες του χρήστη, με ταυτόχρονη παροχή προστασίας της ιδιωτικότητάς του. Για την επίτευξη της εξατομίκευσης της υπηρεσίας με ταυτόχρονη προστασία των προσωπικών δεδομένων του χρήστη γίνεται χρήση ανώνυμων πιστοποιητικών. Γίνεται ενσωμάτωση του συστήματος PrivCAPTCHA στο SIP πρωτόκολλο επικοινωνίας VoIP και πραγματοποιείται πειραματική αξιολόγηση του συστήματος σε πραγματικό VoIP περιβάλλον ανοιχτού κώδικα.

Κεφάλαιο 8: Privacy-preserving Personalised Travel Planning

Περιγράφεται η διαχείριση των προσωπικών δεδομένων των χρηστών εντός του εξατομικευμένου συστήματος προγραμματισμού ταξιδίων myVisitPlanner. Οι χρήστες του συστήματος δημιουργούν προφίλ που περιλαμβάνουν τα χαρακτηριστικά και τις προτιμήσεις τους, όπως και το ιστορικό των ταξιδιών που έχουν δημιουργήσει. Η προστασία των δεδομένων των χρηστών επιτυγχάνεται με τον περιορισμό της πρόσβασης στα δεδομένα από μη-εξουσιοδοτημένους χρήστες, μέσω της χρήσης κατάλληλων τεχνικών κρυπτογραφίας.

Κεφάλαιο 9: Conclusions, Challenges and Directions

Το κεφάλαιο αυτό περιγράφει τις κύριες συνεισφορές της δουλειάς αυτής και παρέχει μια επισκόπηση των ανοιχτών θεμάτων του χώρου και των δυνατοτήτων μελλοντικών εργασιών.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Research Context

Personal information is collected and stored outside the data owners' control during all kinds of online activities today. Apart from computers, the use of mobile applications and sensors lead to constant online presence and subsequent data recording. Sensitive private information is left unprotected at the companies' disposal and is often (un)intentionally leaked to unauthorised parties [9, 66]. This situation can bring great profits to the parties who exploit the information and significant costs to the information owners, both financial and security related.

By giving away their personal information, consumers often do not realise the dangers their actions entail and the possible impact on their lives (identity theft, discrimination, etc). Additionally, even privacy-aware consumers often need to choose between conducting a transaction in a non privacy-preserving way, or not proceeding with the transaction altogether. All kinds of data collected about an individual can potentially be misused. The obvious categories are sensitive information like financial or health records, but the collection of location data or shopping habits should not be considered risk-free either. Electronic trails can be combined to compose an individual's detailed profile. Companies thrive nowadays, making substantial profits by exploiting user personal data.

The economic value of personal information is becoming apparent, as many companies provide their services "for free" and finance themselves through the collection and processing of the personal information of their users. Personal data can be used to achieve more accurate behavioural advertising, more effective service personalisation, as well as better recommendations for products and services. These innovative technologies provide enhanced quality of service for the users, but at a steep price. Users should be able to enjoy the benefits of new technologies without having to give up their privacy. Using privacy-preserving mechanisms there can be a balance between the seemingly contradictory requirements of per-

sonalisation and privacy protection.

At the same time, data accuracy and accountability is an important issue regarding personal data processing. Out-of-date or partial data provide inferior results when it comes to service profiling and personalization. Methods that provide service providers with legitimately acquired, up-to-date and accountable consumer information would bring about important benefits to their efficiency as well.

1.2 Motivation

The availability of new, data driven technologies allows for new, user-centric applications to be implemented. Users are able to receive services tailored to their needs and preferences. This development is definitely useful and contributes to the improved quality of life for individuals, however, at the same time, opens up important privacy and security concerns.

In our view, the necessary technology exists today for better personal information protection to be realised. The ideas and principles already existed and now the enabling technologies are present, both in information science and telecommunications, for them to be put into practice. There are two important elements missing. The first is public demand for privacy protection which will in turn motivate service providers to move towards more privacy-preserving practices. Otherwise, they have little incentive to take on the burden of protecting individuals' personal information [2]. Public awareness should be raised by making it common knowledge that today's lack of privacy protection in online services could be improved. The other missing element is the existence of practical applications, which use the recent technological advancements to create better privacy-preserving conditions for well-meaning companies and individuals. Several applications need to be developed, including a framework for personal data management with an intuitive and user friendly interface, a transactions mechanism for personal information trading, tools for policy editing, licence management and a price estimation mechanism. This work aims at making a few steps towards this direction and our subsequent work will be within this context.

We believe that accountable privacy-preserving services need to be developed to increase consumer trust in everyday transactions and allow for privacy-supporting utilization of the consumer data. Following the principle of least information (or data minimization) [29], customers should be able to perform transactions in a way that only the necessary information about them is revealed to the service provider and preferably remaining anonymous. Service providers on the other side, should be able to utilize the transaction information of their customers to draw useful conclusions about their products, without violating customer pri-

vacy. It would be even better if companies could acquire access to accountable transaction data besides their own to derive general market trends and patterns, compensating the data owners for their service.

The central assumption of our proposed personal data management approach is that personal data can only be stored at the owner's side. One may dismiss this as an unrealistic hypothesis and contend that we cannot count on users to abide by this principle. One may also doubt that there are any incentives to adopt such an approach, especially for the service providers. This proposal may specify an absolute approach for personal data management, it is meant, however, to provide proof of the concept that personal data management can be fair, privacy-respecting and more effective than current practices. What makes our approach possible is that the recent scientific and technological developments and especially the universal acceptance of the Internet have prepared the ground for citizen-centric applications. At the same time, the powerful surveillance and data management tools have contributed to making privacy threats and personal data misuse one of the most important problems in the electronic world.

For the above reasons, we believe that the conditions are mature for investigating alternative paradigms in the management of personal data. The new paradigms should enhance the individuals' control over their personal data. In this context we designed, implemented and evaluated several privacy-preserving applications. The impact and adoption degree of these solutions into practice is a multi-parameter issue, discussed further in the Conclusions (Chapter 9). However, this work constitutes a confirmation that such solutions are feasible.

This work aims at providing consumer-led solutions for personal data management to be used instead of the current company-centric approach. An interesting analogy of the proposed switch in the current practices in the field of personal data management exists with another proposed switch in the field of identity management. According to the Crosby report [40] we should focus on identity assurance instead of identity management:

> At an early stage, we recognised that consumers constitute the common ground between the public and private sectors. And our focus switched from "ID management" to "ID assurance". The expression "ID management" suggests data sharing and database consolidation, concepts which principally serve the interests of the owner of the database, for example the Government or the banks. Whereas we think of "ID assurance" as a consumer-led concept, a process that meets an important consumer need without necessarily providing any spin-off benefits to the owner of any database [40].

We consider it very encouraging that positions about consumer-led solutions are expressed within a very applied context, like the Forum on Identity Management

which prepared the report.

1.3 Methodology

The goals of this research, as already analysed in the motivation, are to develop privacy-preserving applications based on the principle that data owners should be able to control the uses of their personal data, while being able to utilize them to their advantage.

To achieve these goals, applications that need to become privacy-preserving were selected and the following methodology was applied to achieve privacy protection within the application design, while maintaining the intended functionality:

- The properties of the application are determined:
  - The entities that participate in the application.
  - The personal data within the application that need to be protected.
  - The functions that use the personal data.
  - The entities that have access to the personal data.
  - The application output and whether it also contains sensitive information that needs to be protected.
- The appropriate privacy-preserving solution is designed and the appropriate protocols are developed, to achieve the same functionality as the original application, while protecting personal data. Decisions are made on the following issues:
  - Where the personal data are stored in order for the data owners to maintain their control.
  - Where the functionality and the computations of the application take place.
  - Which are the appropriate privacy enhancing technologies that need to be employed to support the privacy-preserving functionality.
- Adversary scenarios are examined and the possible malicious behaviours are identified:
  - The data leaks that may take place in case of malicious behaviour are determined.

  - The possible amendments to the proposed solution to minimize the risk of data leaks are examined.

- Experiments using prototype implementations or real-life applications are conducted to verify the applicability of the proposed solution.

1.4 Results

In the course of this PhD research, the following privacy-preserving applications were created:

1. Polis: A personal data management platform aiming at protecting users personal information during electronic purchases.

   - Protected data: Personal Identifiable Information
   - Joint work. Contribution of this PhD research: Participation in the main research, resulting in the identification of the privacy principles to be followed. Design of the appropriate Polis-based transaction protocols.

2. FPIT (Fair Personal Information Trades): A market for personal information that allows controlled access to individuals' personal data. Recognising the economic aspects of privacy, Polis was further extended by introducing the idea that individuals should be compensated when their personal information is used for commercial purposes, while retaining control of their data. The FPIT work focuses primarily on the economic aspect of the information exchange and the ways the information owners can be compensated when their personal information is being used.

   - Protected data: Personal Identifiable Information
   - Core work of this PhD research.

3. PrivTAM: A privacy-preserving television audience measurement system, incorporating participant compensation functionality. The FPIT idea was applied to create a novel television audience measurement system, which, using SmartTV technology, achieves calculating privacy-preserving, validated ratings, supporting participants' compensation.

   - Protected data: Television viewership records.

- Joint work. Contribution of this PhD research: Application of fair personal information trades techniques for participants' compensation. Participation in the design of the TAM application and its privacy-preserving protocols.

4. Portfolio: A consumer data utilization system for controlled storage and management of user transaction data.

   - Protected data: Personal Identifiable Information and transaction history.
   - Core work of this PhD research.

5. PrivCAPTCHA: A privacy-preserving, user-centric VoIP CAPTCHA system. CAPTCHA challenges used to prevent SPIT calls are adapted to the users' attributes and characteristics, while protecting their privacy. The Portfolio idea is applied to protect user privacy within PrivCAPTCHA.

   - Protected data: User attributes and characteristics.
   - Joint work. Contribution of this PhD research: Application of methods developed in the Portfolio work. Design and development of application-specific anonymous credentials using the Idemix cryptographic library. Integration of anonymous credentials into the SIP environment. Development of appropriate protocols and experimental implementation, using real-world CAPTCHA service and VoIP client.

6. myVisitPlanner: A privacy-preserving personalised recommendation and planing system for tourism.

   - Protected data: User profiles - Characteristics, preferences and travel history.
   - Joint work. Contribution of this PhD research: Design of application and user profiles with privacy protection in mind. Identification of minimum scope for user data. Development of data protection methods within the application processes.

1.5 Structure of the Dissertation

The summarised contents of the dissertation chapters are presented below:

Chapter 2: Background

In this chapter the main concepts regarding the subject of this dissertation are presented. Privacy, security and cryptographic concepts are described as well as the cryptographic building blocks used to construct the solutions proposed in this PhD work.

Chapter 3: The Polis Personal Data Management Framework

The Polis solution for personal data management is described, developed and evaluated. Personal data is stored at the user-side and a user-owned agent controls the access to the data. Representative examples of electronic transactions involving personal data are described and protocols supporting privacy-preserving transactions within Polis are proposed. Experimental results of a prototype implementation and a real-world application are presented, indicating that it is feasible to carry out electronic transactions efficiently, while personal data remain at the owners' side.

Chapter 4: Fair Personal Information Trades

The FPIT (Fair Personal Information Trades) system is described, which proposes a "market" for personal information that allows interested parties to acquire controlled access to individuals' data, in return for fair compensation, either financial or in the form of benefits. Personal information remains at the owner's side and can be accessed under agreed upon conditions via appropriate data licenses. This way, data owners remain in control of their data and are aware of their retrieval and usage purpose. The FPIT architecture utilizes the Polis personal data management framework and incorporates a micropayments system to support large numbers of small amount payments.

Chapter 5: Privacy-preserving Television Audience Measurements

The PrivTAM system proposes a novel, privacy-preserving approach for computing Television Audience Measurement ratings, which utilizes the capabilities of SmartTV technologies. Contemporary cryptographic primitives are utilized to ensure the privacy of the individuals and the ratings' validity. User compensation capabilities are introduced to bring some of the data processing company profits back to the data owners. Experimental results of the Android-based prototype implementation illustrate the feasibility of the approach.

Chapter 6: A Personal Data Management Portfolio for Privacy-preserving Consumer Data Utilization

In this chapter the Portfolio is described, a personal data management system for data owners to store and manage their transaction data, that remains under

their control. Each individual in the system has a Portfolio that is stored locally at their side and contains their transaction history. An encrypted version of the Portfolio data is outsourced to a cloud storage service. Interested parties acquire access to the Portfolio data in a controlled way and always with the user's consent. The proposed solution offers utilization capabilities and simultaneous protection of consumer information, along with accountability assurances for the accessed transaction data.

Chapter 7: Privacy-preserving, User-centric VoIP CAPTCHA Challenges

In this chapter PrivCAPTCHA is described, a system that adapts VoIP CAPTCHA challenges according to the user characteristics, while preserving their privacy. Anonymous credentials are used to achieve privacy-preserving personalisation. PrivCAPTCHA is integrated in the SIP environment and experiments are performed using a real-world, open-source VoIP application.

Chapter 8: Privacy-preserving Personalised Travel Planning

This chapter describes the handling of user personal data within the myVisitPlanner personalised itinerary planning system for tourism. MyVisitPlanner users create profiles that include their interests and characteristics and appropriate activities are recommended. Privacy-protection of the user data is achieved through cryptographic techniques, that limit the data exposure to unauthorised parties.

Chapter 9: Conclusions, Challenges and Directions

This chapter provides the overall conclusions of this work, including the contribution of this PhD research. Additionally, an analysis of the obstacles and directions of the field is provided.

Chapter 2

Background

2.1 Personal Data and Privacy

Personal data is defined as "any information relating to an identified or identifiable natural person ("data subject"); an identifiable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity", according to European privacy regulations [58].

Personal data is in essence any data which can be related to an individual, directly or indirectly. Also referred to as Personal(ly) Identifiable Information (PII) or Sensitive Personal Information (SPI) is information that can be used on its own or combined with other information to identify, contact or locate a data subject.

Privacy is acknowledged as a fundamental human right by Article 8 of the European Convention on Human Rights, defined as a right to respect for ones "private and family life, his home and his correspondence". Information privacy or data privacy refers to the ways personal data can be gathered, stored, processed or disseminated [60].

2.1.1 Fair Information Practices and Privacy Principles

In order to protect personal information, several organisations and countries have issued privacy regulations, which should be followed in order for personal information to be protected; the collectively referred to as Fair Information Practices (FIP).

There are many versions of FIPs, but the core concepts they reflect can be summarised in the following principles [25, 44]:

Consent In order for data processing of an individual's data to be allowed, explicit consent must be given by the data subject with respect to the processing of the data.

Transparency Essentially a prerequisite for consent, as data subjects have to be informed of the purpose, the methods and the recipients of the processing of their data. Additionally, data owners need to be aware of the risks that the processing entails.

Proportionality Personal data may only be gathered and processed in relation to the purpose of the process.

Data minimization Applications should only collect and process the minimal amount of personal information needed for the performed task and should be erased or anonymised as soon as the task is completed.

Information security Addresses confidentiality, integrity and availability, the data protection goals to prevent unauthorised access, manipulation and destruction of personal data.

Accountability The ability to prove compliance with the above principles, i.e. the information gathering and processing was authorised and the data is being properly used and maintained.

2.2 Economics of Privacy

Personal information is collected to construct profiles of individuals, containing information gathered directly and indirectly during their online interactions. Some of the most profitable companies worldwide provide their services to individuals free of charge and make substantial profits by exploiting the gathered information of their users. This fact demonstrates the important economic impact personal information has today.

It is generally acknowledged that personal information has become the "new currency" on the Internet [25, 116], giving even more power to the argument that individuals need to be in control of the dissemination of their information, as well as be compensated for disclosing it [92]. Market mechanisms have been proposed as an appropriate way to address privacy and its economic aspects [139]. The economic aspects of personal privacy are more thoroughly discussed in [139, 4, 11, 88, 92, 89, 76].

An insightful analysis of the different approaches to personal information protection, including market-based approaches is presented in [73]. In that work the conclusion is drawn that it is very difficult to protect from unauthorised data

copying and distribution. This is especially important for personal information, because there is no way of preventing a person allowed to see the data once, from writing it down on a piece of paper. This problem could be addressed by requiring information users to exhibit licences from the information owners, entitling them to use this piece of information for this particular purpose. The benefits from the use of data licences are discussed in [36].

At the same time, data accuracy and accountability is an important issue regarding personal data processing [143]. Out-of-date or partial data provide inferior results when it comes to service profiling and personalisation. Methods that provide service providers with legitimately acquired, up-to-date and accountable consumer information would bring about important benefits to their efficiency as well.

Another interesting point, according to Odlyzko [109], is that privacy violation stems from the need for price differentiation according to the user's profile. This way service providers can determine how much clients are willing to pay for their services. Although price differentiation is usually perceived as a negative idea, it can have positive economic and social results as it allows products and services to be sold in profitable prices (on average) for the service seller and advantageous prices for each individual buyer.

Therefore, the availability of the customers' profile is important for service and price personalisation, as long as privacy protection and data accountability is ensured and the information processing is performed with the individuals' control and consent.

2.3 Privacy Enhancing Technologies and Privacy by Design

Privacy enhancing Technologies (PETs) are the technical mechanisms that play an important role in the quest for practical privacy protection. These mechanisms constitute the building blocks of privacy-preserving systems and services and include encryption, access control mechanisms, anonymous communications protocols, attribute-based credentials and searchable encryption schemes. In [25] the existing cryptographic tools for creating privacy-friendly applications are further analysed.

In order to achieve privacy-preserving systems and services, PETs need to become part of the initial system design and instead of being added later on. This principle, referred to as "privacy by design" is considered essential for addressing privacy requirements in contemporary systems and services, as well as achieving by default compliance with privacy legislation and data protection rules.

Privacy-preserving mechanisms need to be embedded into applications that handle personal data. Anonymisation techniques for published aggregate profile data offer protection while allowing data processing and utilization [14, 103]. These solutions are a positive step toward privacy aware profile data management, however the information owners do not control the handling of their data and do not gain any benefits from the data exploitation. It is important that individuals are protected from unauthorised usage of their information and compensated for the service and profit they offer to legitimate data processing companies whenever their data is processed. Privacy protection should not only be a result of policies and legal measures, rather it should be achieved through a holistic approach that achieves a system's privacy properties as a part of its design [44]. This "privacy by design" principle ensures that the user's privacy is protected in the default system setting, instead of utilizing add-on features to achieve protection.

2.3.1 Private Credentials

Private credentials, based on the Camenisch-Lysyanskaya signature scheme [32, 31] allow users to acquire credentials and demonstrate them without revealing their identity. Using the private credential system described in [33], individuals can use different unlinkable pseudonyms, based on the same credential issued by an identity provider. The private credential system can also provide certified attributes by the identity provider, for the individual to selectively reveal attributes (e.g. their age range, based on their date of birth). Anonymity revocation is also supported by the private credential system, allowing identities of ill-behaving users to be revealed under well-specified conditions. A private credential system, apart from users who receive credentials and the certification authority that verifies user public-private key pairs, comprises of organisations which grant user credentials, verifiers who verify them, and an anonymity revocation manager that allows revocable anonymity.

The ideal system offers unforgeability of credentials, anonymity of users, unlinkability of credential showings and consistency of credentials [31]. Private credentials, also called anonymous credentials, constitute today an accepted and applicable privacy enhancing technology, an open source implementation of which is Idemix[1].

The entities participating in an anonymous credential system are individuals, companies and trusted third parties (e.g. government services), which can assume the roles of issuers, recipients, provers and verifiers. A credential is created by an issuer for a recipient by executing the issuing protocol. The recipient (i.e. the credential owner) can then create a credential proof, to be used by a verifier to

[1]Identity Mixer cryptographic library, https://prime.inf.tu-dresden.de/idemix/

verify the validity of the credential (proving protocol).

To be able to issue anonymous credentials, an issuer needs to generate a public/private key pair and create specifications of the structures of the credentials issued. These specifications and the issuer's public key are then published to be used during the proof protocol.

In order to acquire an anonymous credential, a user chooses a master secret key, according to the agreed upon system parameters (bit length, groups to be used). This secret key enables the creation of multiple unlinkable pseudonyms by the user, to be used with different service providers. The issued credential consists of the issuer's public key, the credential structure (necessary for the verification of its validity) and the attribute values.

During the proving protocol, the prover (i.e. the credential owner) creates a proof on behalf of the verifier that proves ownership of a certain credential. The verifier checks the validity of the given proof. Credential attribute values contained in the proof may or may not be revealed to the prover, according to the settings of the proof creation process.

2.3.2 Searchable Encryption

Searchable encryption solutions provide the ability to store information in an encrypted form, yet allow search operations to be performed on the dataset, in order for authorised users to be able to locate the information they are interested in, without needing to decrypt the whole dataset.

Searchable encryption has been a very active research topic in the past decade, yielding increasingly efficient and applicable solutions. Applying technologies from cryptography, data structures, algorithms, information retrieval, and databases, the state-of-the-art constructions achieve different trade-offs between security, efficiency, and query expressiveness [84]. Kamara presents in [84] an overview of the research area, the different solution approaches and the challenges left to address. A thorough and comprehensive survey on searchable encryption schemes and the characteristics of the different approaches by Bösch et al. can be found in [23].

Searchable Symmetric Encryption (SSE) is a practical method for searchable encryption that provides a reasonable trade-off between efficiency, functionality and security. A *searchable symmetric encryption scheme* consists of an algorithm EDBSetup and a protocol Search fitting the following syntax [80]: Let λ be a security parameter. A database $\mathrm{DB} = (\mathrm{ind_i}, \mathrm{W_i})_{\mathrm{i=1}}^{\mathrm{d}}$ is a list of identifier and keyword-set pairs, where $\mathrm{ind_i} \in \{0,1\}^{\lambda}$ is a document identifier and $\mathrm{W_i} \subseteq \{0,1\}^*$ is a list of keywords in that document. EDBSetup takes as input a database DB and a list of document decryption keys RDK, and outputs a secret key K along with an encrypted database EDB. The search protocol proceeds between a client

C and server E, where C takes as input the secret key K and a query (a tuple of keywords and a boolean formula) and E takes as input EDB. At the end of the protocol C outputs a set of (ind, rdk) pairs while E has no output.

The dynamic multi-client OXT scheme consists of the basic OXT scheme [35], appropriately enhanced to support multi-client search functionality [80], and dynamic updateable encrypted datasets [34]. Additionally, the OXT scheme supports conjunctive search and general boolean queries and remains efficient for very large databases. The practical applicability of the scheme is demonstrated by experimental results both of the OXT prototype implementation and in real-world applications [34].

In the basic OXT protocol the server holds encrypted pointers to documents in a dictionary D and the search client (and data owner) holds a list of keywords. The protocol output is the set of encrypted pointers to the documents containing the client's keywords. To retrieve the documents, the client decrypts the pointers and obtains the matching (encrypted) documents. The server does not perform the decryption and cannot learn the keywords in the client's query. The OXT scheme computational complexity is independent of the number of documents in the dataset and scales with the number of documents matching the least frequent keyword in the search query.

To support multi-client search functionality, the basic OXT scheme is enhanced so that the data owner outsources an encrypted dataset to an external server and allows other parties to perform queries on the encrypted data by providing them with search tokens for specific queries. In order to achieve that, the data owner provides the search-client with a set of trapdoors (determined by the query and independent of the searched data), that can be transformed into search tokens for the OXT scheme. To ensure that the searcher is authorised by the data owner, the trapdoors provided by the data owner are homomorphically signed and so are the transformed search tokens submitted to the server. These tokens and signatures are unforgeable even by fully malicious clients [80]. Homomorphic signatures enable a signer to create a signature on a document and allow other parties to make predefined alterations to it and obtain a new signature on the altered document without interaction with the signer [82, 13].

To support dynamic additions to the encrypted dataset after it has been uploaded, the dataset owner needs to be able to compute the labels for the new data, to be added to the dictionary by the server. The scheme remains the same as the static OXT scheme, with the addition of an Update protocol and an auxiliary encrypted database EDB^+ and dictionary D^+ are used, which are initially empty and change as updates happen. Searches in the dynamic databases are performed by the server by first searching the dictionary D for keyword w, as in the static case, then re-computing all the labels corresponding to w in D^+. For this opera-

tion the data owner needs to provide the value of the keyword-specific counter, therefore a dictionary D_{count} is maintained, associating each added keyword with its counter value.

2.4 Related work

Some important research projects on personal data management that also advocate using cryptographic protocols to develop privacy-preserving applications and identity management are Prime [12], PrimeLife [118], ABC4trust [125] and FutureID [124]. These projects have yielded important results on the topic of privacy-preserving mechanisms and usable privacy, including the extension and improvement of Idemix[1], an anonymous credential system and library, used as a building block in several of our proposed solutions.

Prime's approach is that data is provided to data consumers for them to store under well-defined policies. Within Prime, advanced techniques for privacy policy negotiations and enforcement were designed and implemented. The PrimeLife project [118] is the successor of Prime, examining lifelong privacy maintenance and privacy protection in emerging Internet applications. An innovative product on personal data management that caught the attention of major companies, such as IBM and Microsoft, is U-Prove [134].

There is a number of interesting works that utilize cryptographic primitives to achieve accountable privacy-preserving services and maintain the notion that the information owners should retain control of their data. A privacy-preserving National Identity Card is proposed in [48], where a smart card is used to store individuals' personal information and allows them to prove binary statements about themselves. A project that focuses on personal data of smart phones is the openPDS personal data store [47, 46]. The openPDS system supports a mechanism for installing third-party apps which can have controlled access to the personal data of the data store. In [52], Pythia, a mobile contextual suggestion system for tourism is presented. Pythia abides by the principle that all data of the user reside on the user-side and are not disclosed to any other party.

[1] Identity Mixer, http://www.zurich.ibm.com/idemix/details.html

Chapter 3

The Polis Personal Data Management Framework

3.1 Introduction

In order to enhance privacy protection during electronic transactions, we propose a personal data management framework called Polis, which abides by the principle that individuals maintain control over their personal data, which reside only at their own side. We assert that electronic transactions can be feasible, whilst personal data resides at the individuals' side. To support this claim, we design, build and evaluate the prototype system Polis [55, 56], which implements the above principle. We identify representative electronic transactions that involve personal data and propose Polis-based protocols for them. We show that Polis can satisfy important data protection principles in a natural and efficient way and describe how Polis can be integrated into online transactions to manage personal data. The results of this work indicate that electronic transactions can remain both feasible and straightforward, while personal data remain only at the owner's side.

3.2 Related Work

The idea that individuals should retain ownership of their personal information themselves and decide how this information is used, is discussed in [96]. A point made in [127] is that, although considering personal data the owner's private property is a very appealing idea, it would be rather difficult to practically apply it and legally enforce it. Our approach proposes an idea that has the same practical effect as considering personal data the owner's private property, but withdraws the legal objections involved with this idea. The argument that per-

sonal data would be safer at the user's side is also examined in [106].

Different kinds of frameworks that are related to personal data have recently been proposed or are in progress. In particular, privacy sensitive management of personal data in ubiquitous computing is discussed in [77], storing personal data in an individual's mobile device is examined in [79]. Protecting personal data that is stored within a company is considered in [126, 86]. More related to Polis is a rich but also complicated framework for privacy protection, proposed in [102]. This framework is built on the principle that personal data is kept inside a "Discreet Box", located at the service provider's side. An agent-based solution to address usability issues related to P3P (Platform for Privacy Preferences Project) is presented in [99]. Other results in this field, less related to Polis, can be found in [65, 17, 75]. General surveys on privacy enhancing technologies are given in [66, 76, 74]. To our knowledge, Polis is the first general framework for managing personal data only at the owner's side.

3.3 The Polis Approach

The Polis approach is based on the following principle:

"Polis-users are prohibited from storing any personal data but their own."

Polis is meant to be employed by privacy concerned internet users which fulfill the requirements of having:

- A reliable, always-on access to the Internet, in order for their agents to be always accessible.
- A certificate from an approved Certification Authority.

We design, implement and evaluate a Polis prototype and show that the above simple and straightforward assumptions suffice to build a personal data management framework that works seamlessly with online transactions. The Polis prototype and its evaluation are described in Section 3.6.

3.3.1 Polis Concepts and Architecture

At this point we consider it necessary to introduce a few terms that will be used in this work:

- In Polis, personal data refers to primitive personal information of individuals like name, birth date, address, etc. Personal data corresponds to what is called *off-line identity* in [2]. Our focus is on privacy-enhanced management of the off-line identity.

- An individual Internet user is a potential customer who can purchase either goods or services. This user can be called *individual*, *customer* or *data subject*. We will use the terms *individual* and *customer*, interchangeably.
- An entity that provides the aforementioned goods or services can be called *shop*, *company*, *service provider* or *data controller*. We will use the terms *shop*, *company* and *service providers*.
- Both individuals and companies can become *Polis-users*.

Every Polis-user is represented by a dedicated entity. This entity can be used to instantiate a corresponding Polis-agent, which is the main architectural component of Polis. The agent is used to manage the personal data of the entity and to provide controlled access to it. Service providers use the agent to retrieve personal data from affiliated users. The general architecture of Polis, as well as the constituents of a customer agent and a shop agent are presented in Figure 3.1.

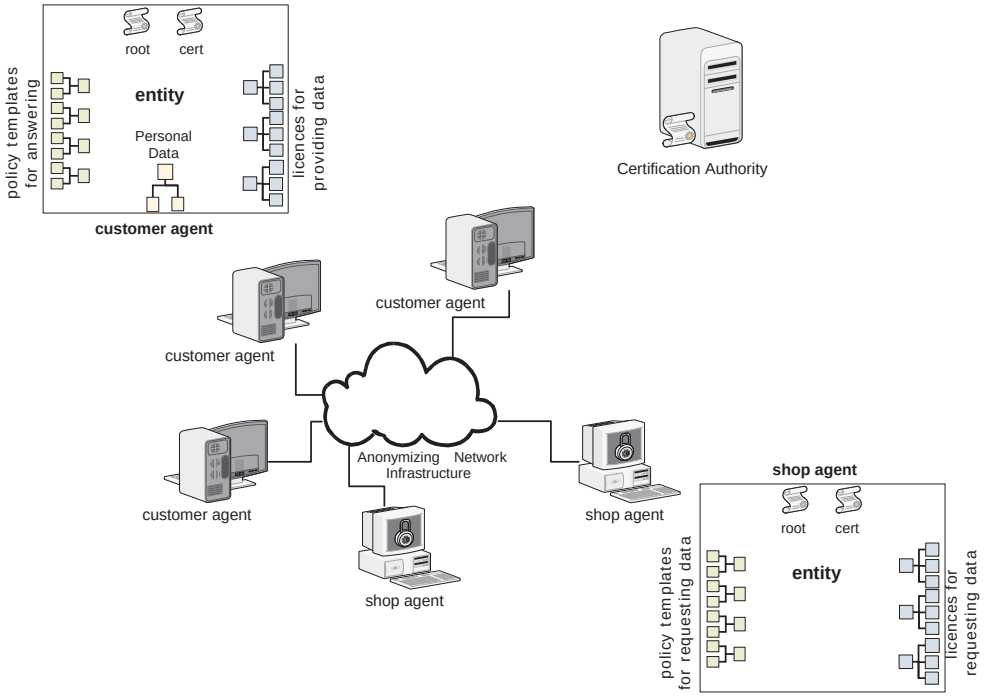

Figure 3.1: The Polis architecture.

We would like to emphasise the following characteristics of the Polis architecture:

- From the service provider's point of view, Polis provides a decentralised approach for the storage and management of personal data.
- On the contrary, from the customer's point of view, Polis is a fully centralised system, in the sense that personal data is located and managed locally by the owner's agent.

3.3.2 Schemes for Personal Data and Policies

Critical components for a personal data management framework like Polis are the schemes for representing personal data and policies. Some known schemes for personal data are P3P [142] and CPExchange [22]. Approaches for policies related to personal data are also discussed in [85, 86], while related work on personal data and policy schemes can be found in [51].

We use schemes that are simple, yet powerful enough, for the needs of the Polis prototype. Examples of a personal data scheme and a policy, as used in Polis, are shown in Figure 3.2.

Policy

```
<User Enabled="true" Entity=" eshop">
 <Name>
  <Given>
    <Permissions>
      <License Purpose="shipping">
        <GrantAccess>true</GrantAccess>
        <DateTime>
          <Start>2008-01-01 00:00:00</Start>
          <End>2008-12-31 23:59:59</End>
        </DateTime>
      </License>
      <License Purpose="billing">
        <Count>3</Count>
        <DateTime>
          <Start>2008-01-01 00:00:00</Start>
          <End>2008-12-31 23:59:59</End>
        </DateTime>
      </License>
        ...
    </Permissions>
  </Given>
 </Name>
</User>
```

Personal Data

```
<User>
 <Name>
   <Given>John</Given>
   <Family>Doe</Family>
 </Name>
 <Home-Info>
   <Postal>
     <Street>Nowhere Street 001</Street>
     <City>Deadend</City>
     <StateProv>Ouitcy</StateProv>
     <PostalCode>11111</PostalCode>
     <Organization>DIPH</Organization>
     <Country>Neverland</Country>
   </Postal>
   <Telecom>
      ...
   </Telecom>
 </Home-Info>
</User>
```

Figure 3.2: Examples of a personal data scheme and a policy.

There are eight general categories of personal data in Polis, organised hierarchically, namely Name, BDate, Cert, Skill, Characteristic, Home-Info, Business-Info and CreditCard. Each of them has one or more subcategories. The terminology

used is based on P3P for the user information part, with the addition of the financial information (CreditCard) taken from CPExchange, plus the extra personal information fields (Skill and Characteristic). Each entity stores its personal data in a local XML document.

The components of a policy are the following:

- *Principals*: The Polis-entities.
- *Data*: Every single item of a user's (Polis-entity) personal data.
- *Purposes*: The set of purposes that entitle principals to retrieve data.
- *Usage restrictions*: Additional restrictions exist that limit access rights to a specific number of accesses or a specific time interval, or both.

Other important concepts of Polis are the licence and the contract. A *licence* comprises of the data involved, the valid purposes that allow data retrieval, as well as the rules to provide either full or restricted access. The use of licences to protect personal data is discussed in [36, 93, 59]. A *contract* concerns two principals and an arbitrary set of licences. An agent can sign any number of contracts with an arbitrary number of entities.

3.4 Incentives and Objections

In this Section we discuss incentives and objections for the Polis approach and provide arguments that the Polis solution can be beneficial not only to individuals but also to (well meaning) shops.

The fact that a Polis-user's personal data must be retrieved from the owner's side every time it is needed, automatically fulfills many critical requirements found in FIP-like regulations. Moreover, the way the Polis framework can be integrated into database management systems automatically fulfills the requirements of Hippocratic databases [7]. Besides individuals, shops can also obtain important benefits from the adoption of Polis, like a more privacy-friendly profile, simplified data maintenance and data cleansing, as well as significantly reduced responsibility for the safety of customers' personal data.

3.4.1 Incentives for Individuals

1. *Polis-users maintain maximum control over their personal data.* They are able to monitor, at any point in time, all the contracts they have signed, as well as the time, purpose and principal of each data item access that has taken place. Unlike the Polis approach, current practices result in inability

to keep track of where one's personal data reside and how often they are being accessed.

2. *Individuals can trivially exercise their right to up-to-date personal data.* A user simply has to update the locally stored record. Each time a company wants to access it, it will be retrieved on-the-fly from the person's agent and not from the company's outdated database. Consider the act of updating one's address or telephone number. With current practices, the individual has to recall every peer that rightfully possesses this information and go through a record update process for each such peer.

3. *Individuals can handle all kinds of privacy-related rights and preferences through their Polis-agent.* To this end, a unified user interface is being used, while personal data disclosure takes place through a clear data flow. These attributes are absolute requirements for the effective privacy-enhanced management of personal data [98].

4. *The risk that the privacy of the individual is violated due to data breaches from company databases is significantly reduced.* Just like credit cards, Polis-contracts can be cancelled to become useless to invaders. Even if a company does not realise that its database has been compromised, invaders will have to acquire the company's private key, in order to be able to use the stolen contracts. Even in that case, the invaders will only have access to the particular data that the contracts authorise this company for. Furthermore, the data owner will be able to know which data items were leaked and when this happened.

5. *Privacy-concerned individuals will no longer have to choose between either giving away their personal data or not conducting an electronic transaction.* Nowadays individuals suffer the coercion that occurs when there is only one reasonable way for them to receive certain needed services or information [105], i.e. by giving away their personal data. Furthermore, according to Acquisti in [2]: "... as merchants decide against offering anonymizing technologies to their customers, the privacy concerned customers choose not to purchase on-line, or to purchase less. A latent, potentially large market demand remains therefore unsatisfied". We believe that approaches like Polis can offer a viable alternative to current practices for personal data management.

3.4.2 Incentives for Service Providers

1. *The customer's personal data in the shop's database remains always up-to-date.* In addition, this is accomplished without any maintenance costs for the shop.

2. *The use of Polis contributes to improved data quality and can simplify the data cleansing task.* Data cleansing is the act of detecting and removing and/or correcting a database's dirty data (i.e. data that is incorrect, out-of-date, redundant, incomplete, or formatted incorrectly). Data quality is a critical factor for the success of enterprise intelligence initiatives and can incur costs and delays to company operations [119, 140].

3. *Polis releases the shop from the burden and responsibility of keeping customer data safe.* The shop is freed from a set of serious responsibilities for protecting customers' personal data and the risk of being considered liable for serious data breaches. Incidents of intentional or unintentional data breaches are unfortunately quite common and a reasonable worry is that a lot of them never reach the attention of the media. Some representative examples of such situations are the Choicepoint case, a data broker who sold private records of over 150,000 Americans to a group of criminals in 2005 [39], the incident that took place in the UK, where two computer discs containing the personal data of 25 million citizens were lost in the post [9], as well as the Deutsche Telecom incidents [121].

4. *Polis promotes a more privacy-friendly image for the service providers that adopt it.* The commitment that the shop does not store any personal data locally is an appealing argument for privacy-sensitive customers.

5. *Polis can be integrated into a company's existing information system.* As we illustrate in Section 3.6, Polis can naturally handle heterogeneous sets of customers, consisting of both Polis and conventional ones. This fact removes the counterincentive of companies having to go through a demanding transition process in order to integrate Polis into their systems. The company does not get tied down by Polis into having only Polis users in its database.

3.4.3 Potential Objections for Individuals

1. *Managing a personal agent is by definition a critical task, prone to errors and omissions by the user.* However, being in charge of the data disclosure process through a unified procedure like Polis, is much more convenient and

effective compared to current practice, as described in the third incentive for individuals.

2. *Considerations about the agent's security.* An individual's Polis-agent contains critical personal data and digital agreements for data access. Consequently, a production-ready Polis-agent should satisfy high security levels. We believe that this is a viable task, since the Polis-agent has a precise, well-defined functionality and can be operated behind firewalls on a user-controlled computing platform. Moreover, the decentralised approach of Polis for personal data can also contribute to improved data security, since invaders find large collections of personal data much more inviting than an individual's personal data [106].

3. *Polis does not protect individuals from malicious shops that misuse personal data.* Nevertheless, a malicious shop in Polis cannot cause more damage than it could cause with current practices.

3.4.4 Potential Objections for Service Providers

1. *Losing control of customer data.* This objection does not really apply to the Polis approach since service providers will still have access to the data they are entitled to. Well-meaning parties will not lose control over their customer's data. Internet connection reliability could also be an issue for Polis, but as already mentioned, it is widely accepted that reliable Internet connectivity can be considered a given nowadays and in the future. Besides, Polis does not restrain companies from keeping records of customer's profiles. These records will not contain any data of the customer's offline identity and will resemble pseudonymous data processing.

2. *The adoption of Polis can cause significant overheads to company processes.* The possible delays in data retrieval caused by the employment of Polis should not be a hindering factor for its adoption. Retrieval of personal data is neither a task that is carried out frequently, nor a time critical process, therefore these delays will not affect the efficiency of the company procedures.

3. *Service providers could be scammed from malicious Polis-users.* Polis-contracts and licences constitute proof that a service provider has the right to access the specified Polis-user data. Therefore, when needed, a company can resort to the appropriate actions. The CA or some other designated trusted third party could be used to settle such cases.

3.4.5 Enforcement and Detection

An important aspect of every (electronic) contract is the ability to verify and enforce that the parties will not violate its terms. Polis can handle detectable privacy breaches, i.e. breaches for which data released to the shop finds its way back to the individual who submitted that information [71]. In this case a Polis compliant shop must be able to present evidence that those data were rightfully obtained for the specific purpose, at the specific time, using data licences [36]. A more challenging task would be to detect Polis-shops that leak customer's personal information. A relevant problem is discussed in [71].

Due to the very nature of personal data, it seems that once a service provider possesses some data, there is no technically feasible way for absolute abuse prevention. Consequently, apart from technical measures, we will have to rely on market, legal and social dynamics for handling personal data properly ([71, 6] and [77, Section 5.8.5]).

As far as violations from the user's side are concerned, if the terms of an agreement are violated and the individual refuses to fulfill the contract-defined obligation of providing personal information, then the service provider can use the customer-signed licence to prove entitlement to access the data.

3.5 Polis Applications

In this Section, we discuss how Polis can be used within common electronic transactions and present indicative higher level applications that can be built on top of a decentralised personal data management framework like Polis.

3.5.1 Polis in Common Transactions

Polis can be potentially employed in any transaction where users have to submit (some of) their personal data. The overall process is outlined below:

Polis in transactions. When users have to fill in forms with their personal data, they instead provide the contact details of their agents. The agents of the shops and the users/customers establish agreements. A successful agreement grants access to the customer's private data for the specific data items and the amount of time needed to complete the transaction. In Figure 3.3 the interaction of Polis-entities within a Polis-transaction with an e-shop is illustrated.

This process can be used for registrations at e-shops, portals and other online services. In general, any application that involves personal data, like identity

management systems [117] and e-government platforms can be supported. The need for privacy protection in e-business applications is stressed in [89]. The ease of employing Polis lies in the fact that it can work as a middleware, which takes care of the personal data exchange between parties in higher level applications.

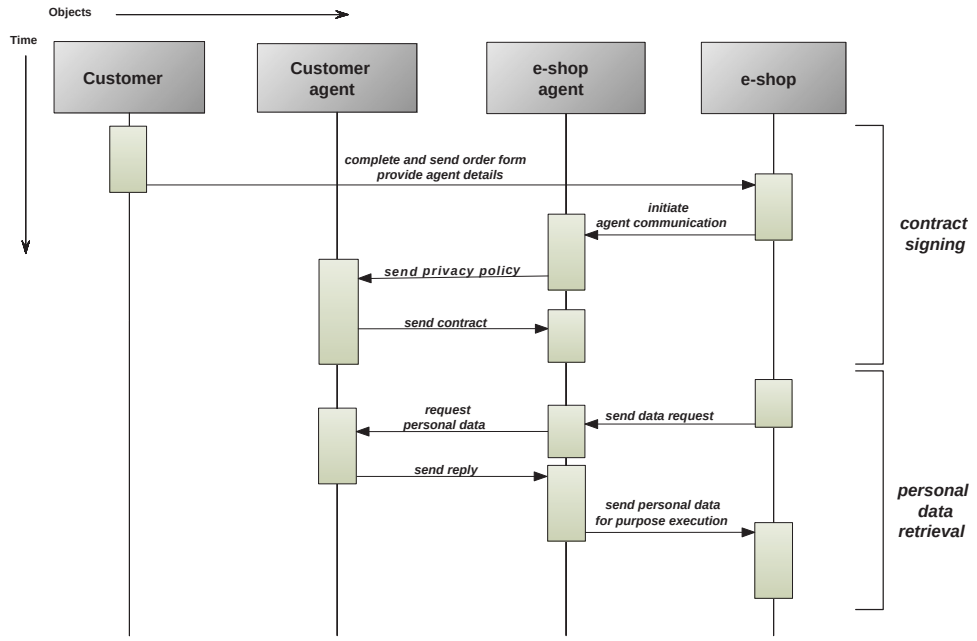

Figure 3.3: Polis in transactions with an e-shop.

3.5.2 Prospective Applications for Polis

An infrastructure like Polis can be a realistic step in the direction of effectively controlling personal information. Apart from the direct gains of using Polis in every day electronic transactions, there are some interesting possibilities for higher level applications that could utilize it.

Microtrades and Information Markets. Polis could be utilized to facilitate personal data exchange in personal-level microtrades between Polis-agents. Such an application is examined in [131]. Polis-users can give permission to information gathering companies to access (some of) their personal data, for an agreed price. Each time a company needs to regain access to them, the agreed amount of money should be paid. Furthermore, Polis could provide the ground for more advanced financial applications for personal data. The market for personal data described by Laudon in [96] is an example

of such applications. In particular, Laudon proposes the so called National Information Market (NIM), where personal information can be traded in a National Information Exchange. The adoption of a framework like Polis would simplify the evolution of NIM-like infrastructures.

Privacy-enhanced ubiquitous computing. Online data of individuals can be conveyed through their Polis-agents. In this case, Polis could work as an open architecture for ubiquitous computing applications. For example, dynamic location information could be retrieved from the individual's Polis-agent, like the rest of the individual's personal data.

3.6 The Polis Prototype

We designed and implemented a Polis prototype with the main objective to demonstrate that electronic transactions are feasible while personal data remain only at the owner's side. Another technical objective of the development of the Polis prototype was to make its deployment simple and friendly to contemporary information management practices. We believe that we have fulfilled the above goals adequately. Furthermore, a fully developed Polis platform should be able to satisfy the general properties that privacy-enhancing technologies must hold in order to be considered useful, according to [66].

3.6.1 Technologies of the Prototype

The basic technologies used to develop and employ the Polis prototype are:

- The Eclipse IDE and the Java programming language to create portable, platform independent tools.
- A Pubic Key Infrastructure (PKI) is used for creating trusted certificates according to the X.509 standard. For demonstration purposes, an elementary Polis CA has been developed to be used in experiments. In a real world application, a commercial CA could be utilized.
- User data, policies and contracts are represented as XML documents.
- The Tor anonymizing infrastructure is used optionally to achieve anonymity for the clients and/or implement agents as hidden services [50].
- The Derby embedded database server is employed internally by the agent for its data storage needs.

- Bouncy Castle's security provider is used for cryptographic primitives.
- The database case study has been implemented on an Oracle database management server (DBMS). Similar integrations of Polis should be feasible with other popular DBMSs like IBM DB2 or Microsoft SQL Server.

3.6.2 Deploying the Polis Prototype

In order to deploy Polis:

- Customers install the Polis-agent, store their personal information and prepare the necessary policy templates.
- Companies install the Polis-agent, prepare policy templates and integrate the agent with the company's information system. *Polis-customers can co-exist with normal customers at the company side.*

3.6.3 Polis Collaborating with a Database Management System

Polis can be incorporated into the back-office of a company and take care of the personal data management. This is accomplished by integrating Polis with the company's database management system. The basic idea is that personal data fields do not contain the actual data, instead, a ticket (represented by an appropriate object) is used to retrieve the data value on the fly. We tested Polis with an Oracle database server. The approach is illustrated in Figure 3.4.

The integration was straightforward. Two Java Stored Procedures (JSP's) and a small set of triggers and database views were sufficient to implement the connectivity between Polis and the database server. It is noteworthy that, using simple object-relational features, as well as views and triggers, the Polis enhanced database can be operated as a normal one, while the Polis related operations are transparent to the database user.

3.6.4 Experimental Evaluation

We prepared an elementary Polis environment with a set of Polis-agents installed on the local network of our laboratory. A snapshot of a Polis-agent's GUI is shown in Figure 3.5. A set of web pages, including web forms and dynamic web pages, were used to support experiments. The customer database contained 27 customers in total; 11 conventional customers and 16 Polis-customers (4 of which used Tor hidden services [50] for their agents). We performed an extensive set of experiments within the above Polis environment. The experiments

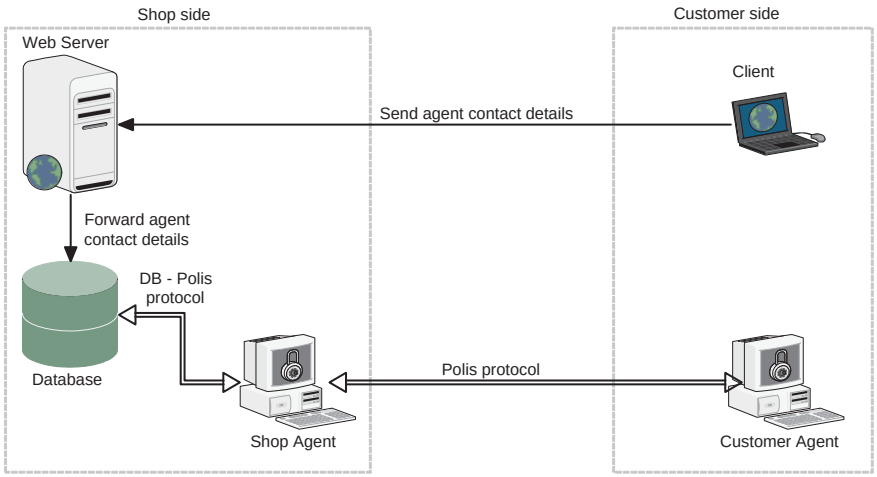

Figure 3.4: An interaction example of Polis-entities. A shop uses a Polis-enabled database for customer registration and personal data retrieval.

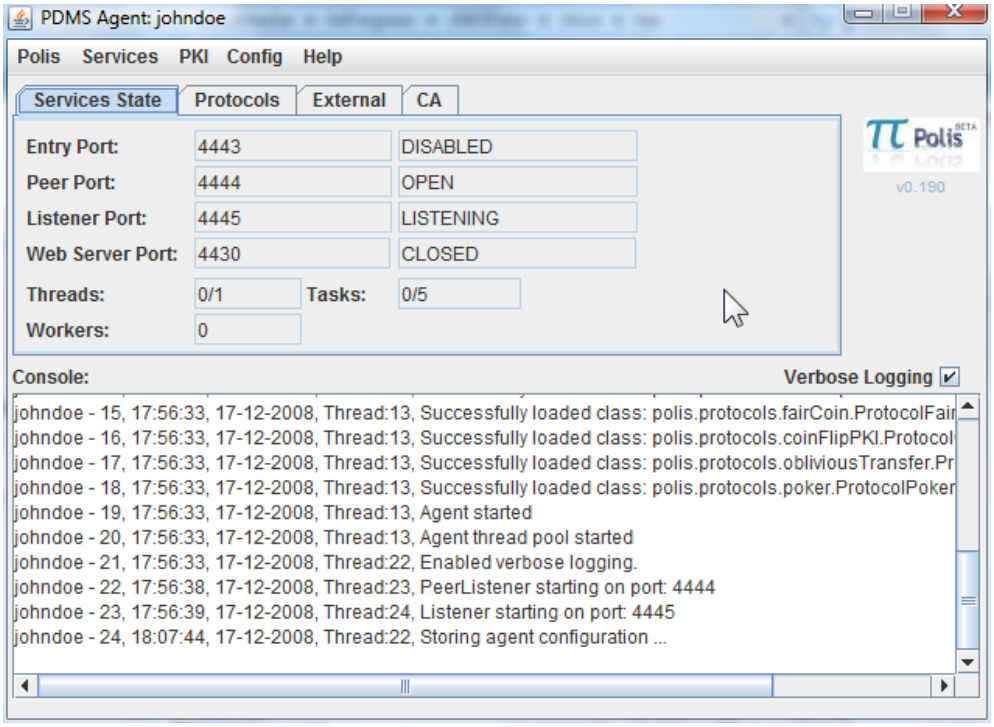

Figure 3.5: A Polis-agent's GUI snapshot.

involved database operations on tables with Polis data to verify their integration into the database. In particular, we executed some representative *insert* and *select* operations on the customer table, a *join* operation between two tables and

created some *views* in the database. Both Tor-enabled Polis-agents (the agent is accessed through a Tor hidden service) and conventional Polis agents were tested. As expected, all the operations were accomplished successfully. Moreover, the Tor-enabled agents operated indistinguishably from the conventional agents with the exception of occasional timeouts, due to the Tor network itself.

ID	First Name	Last Name	Street	City	Postal Code	Country	State	Organization	User Type
23	Sayid	Jarrah	Goodmind 125	Kenya	87484	Kenya	Kenya	TN	**Non-polis User**
22	Walter	Skinner	Filosofou 145	Florina	No Permission	No Permission	No Permission	EU	**Polis user**
21	Fox	Mulder	Kolokotroni 40	Athens	10156	No Permission	No Permission	No Permission	**Polis user**
20	Joey	Tribbiani	Central Perk 46	New York	74889	USA	New York	US	**Non-polis User**
19	Dana	Scully	Paranormal Street 125	Deadend	No Permission	No Permission	No Permission	No Permission	**Polis user**
18	Phoebe	Buffay	Central Perk 23	New York	47952	USA	New York	US	**Non-polis User**
17	Veronica	Donovan	Xonoloulou 210	Tokyo	No Permission	No Permission	No Permission	No Permission	**Polis user**
16	Alexander	Mahone	Fox River 15	Washington	8769	USA	Washington	US	**Non-polis User**

Figure 3.6: Report from the customer table of a Polis-enabled database. The table contains both, Polis and non-Polis, users.

3.6.5 A First Case Study

We performed a preliminary case study on the integration of Polis with a content management system (CMS). More precisely, we integrated Polis with the Elxis CMS, an open source CMS released under the GNU/GPL licence. We selected Elxis for our case study because it is a fully functional CMS, it is open source and it supports the Oracle DBMS. A number of extensions exist for Elxis that enrich its functionality; one of them, the IOS eshop component, turns Elxis into an e-shop.

The integration process was straightforward. A first version of a Polis-enabled Elxis CMS instance was working in beta status (Figure 3.7) after less than a man-week's work. Figure 3.8 shows how the profile of a specific user appears in the Elxis application, at different time instances. Auditing the user's Polis-agent reveals each access to the user's personal data.

3.6.6 Evaluation Conclusions

The evaluation of the Polis prototype and the Polis case study proved the feasibility of the main Polis approach and confirmed the features of the Polis approach discussed in Section 3.4. A comparison of the features that are available to Polis-enabled individuals in contrast to conventional/non Polis-enabled individuals of the Elxis CMS-based application is shown in Table 3.1. The table compares the Polis approach to the current practice in personal data management. To this

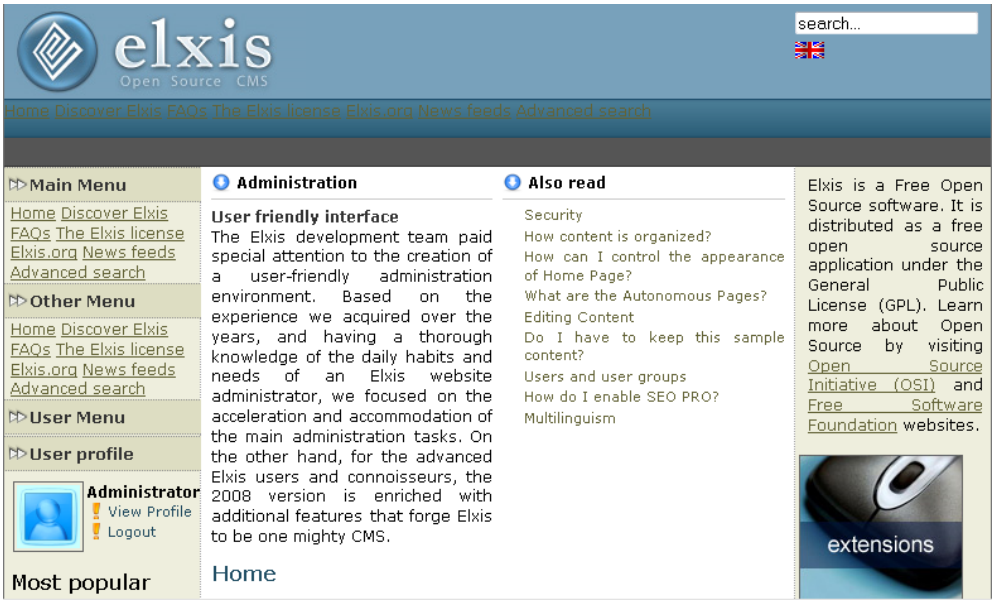

Figure 3.7: The Polis-enabled Elxis CMS instance.

(a) The Elxis user profile.

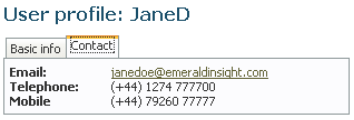

(b) The profile after the user changed his e-mail address.

(c) The profile after the data licence for the telephone field has expired.

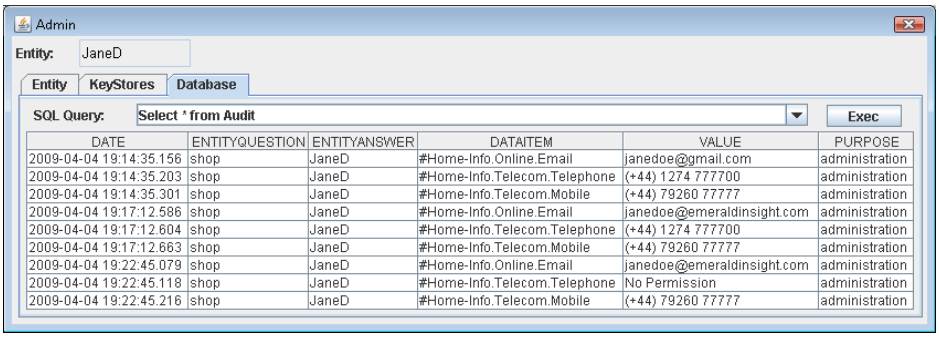

DATE	ENTITYQUESTION	ENTITYANSWER	DATAITEM	VALUE	PURPOSE
2009-04-04 19:14:35.156	shop	JaneD	#Home-Info.Online.Email	janedoe@gmail.com	administration
2009-04-04 19:14:35.203	shop	JaneD	#Home-Info.Telecom.Telephone	(+44) 1274 777700	administration
2009-04-04 19:14:35.301	shop	JaneD	#Home-Info.Telecom.Mobile	(+44) 79260 77777	administration
2009-04-04 19:17:12.586	shop	JaneD	#Home-Info.Online.Email	janedoe@emeraldinsight.com	administration
2009-04-04 19:17:12.604	shop	JaneD	#Home-Info.Telecom.Telephone	(+44) 1274 777700	administration
2009-04-04 19:17:12.663	shop	JaneD	#Home-Info.Telecom.Mobile	(+44) 79260 77777	administration
2009-04-04 19:22:45.079	shop	JaneD	#Home-Info.Online.Email	janedoe@emeraldinsight.com	administration
2009-04-04 19:22:45.118	shop	JaneD	#Home-Info.Telecom.Telephone	No Permission	administration
2009-04-04 19:22:45.216	shop	JaneD	#Home-Info.Telecom.Mobile	(+44) 79260 77777	administration

(d) The user's Polis-agent audit.

Figure 3.8: The profile of a Polis-enabled Elxis user at different time instances and the corresponding entries in the user's Polis-agent audit.

end it highlights a set of important advantages/disadvantages for Polis-users of the Polis-enabled Elxis CMS. The comparison should be valid for a wide range

of possible Polis-enabled e-business applications.

		Polis-enabled user(s)	Conventional user(s)
1	**awareness**	the individual is aware of any access to his data	the individual receives no information
2	**data opt-out**	trivial	the individual has to contact the Elxis shop
3	**specifying policies**	the individual can specify his own data access policies or adopt proven policy templates	the individual has to rely on the company specified privacy policy
4	**security**	the individual is not affected by most types of attacks on the Elxis shop	any data leakage from the Elxis shop may affect the individual
5	**effort**	the individual has to manage his own data	the individual simply gives away his data
6	**delay**	the data is retrieved from the Polis-agent of the individual	data is retrieved from the company database
7	**data cleansing**	trivial (the data is retrieved from the Polis-agent)	data entry errors may occur
8	**data update**	trivial (happens implicitly)	the individual has to go through a proprietary procedure

Table 3.1: Advantages/disadvantages for Polis-users of the Polis-enabled Elxis CMS instance.

The evaluation of Polis also revealed important improvements that are possible for the Polis-agent and the accompanying tolls. One improvement concerns the usability of Polis. In the relevant literature it is pointed out that the use of an appropriate Graphical User Interface (GUI), that clearly demonstrates concepts for expressing privacy preferences, is very supportive for effective use of privacy protection systems [98, 112, 81, 1]. A first tool that we have created is a Polis Add-on for Firefox, which, amongst other features, alerts the user each time personal data is requested (Figure 3.9).

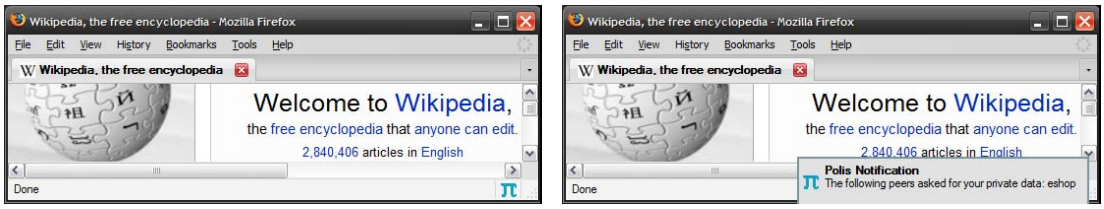

(a) Internet browsing moment of a Polis-user. (b) The user's data have been requested.

Figure 3.9: The Polis Add-on for Firefox alerts the user when personal data is requested.

3.7 Discussion

In this work, we design, implement and evaluate Polis, a personal data management framework which embodies a fundamental privacy principle: Personal

data of individuals reside only at their side. Polis aims at making storage of personal data unnecessary for contemporary online transactions to work efficiently. This way, users will be able to monitor and limit the distribution of their personal data, according to their needs and preferences. Furthermore, the safety of stored personal data is enhanced and personal data accuracy is ensured.

In conclusion, this work demonstrates the fact that it is possible to deploy a privacy-enhancing prototype like Polis, in order to achieve significant privacy protection, in the current electronic world. We cannot expect Polis to become a panacea for all kinds of privacy problems. However, we believe that Polis has more advantages than disadvantages compared to current practices for personal data management. Finally, it is very encouraging that given one basic assumption, the transition to a personal data protecting way of conducting online transactions, can be natural and smooth.

Chapter 4

Economics of Personal Data Management: Fair Personal Information Trades

4.1 Introduction

Individuals today have no control over the way their personal information is being used even though they are the ones to suffer the consequences of any unwanted uses of their information. We propose addressing this externality through the creation of a market for personal information, where licenses to access individuals' personal information will be voluntarily traded. Through this market, satisfactory compensation to the information owner is provided, whilst personal information remains under the owner's control. Using cryptographic tools and micropayments we propose and develop a prototype for personal information trades where the above principles are implemented and tested.

We consider the use of the Personal Information Market (PIM) concept, where access to personal information can be legitimately exchanged, providing at the same time mutual benefits to companies and individuals. In order for a PIM to be effective many challenges need to be addressed. First, it is important that personal information is exchanged in a way that misuse attempts are prevented or deterred. In this work we address this challenge by using data licenses, which are based on cryptographic primitives. These licenses provide access to individuals' personal information under well specified conditions and therefore ensure that information owners do not lose control over their information. Another challenge to be addressed in realizing a PIM is the payment scheme. Information owners ought to be compensated for providing access to their personal information. The amount, however, of compensation per information item access should be small.

Therefore, common bank/credit card payments are not suitable for this kind of transactions, since they are relatively costly and time consuming. It should be possible to conduct large amounts of transactions efficiently and without entailing added costs. An appropriate payment method for this PIM is the idea of micropayments [123], which allows for small amounts of money to be exchanged, in an easy, efficient and inexpensive way.

We believe that an interesting analogy can be drawn between modern Internet users' privacy rights and developing countries producers' rights. Just like producers in such countries are vulnerable to work exploitation and suffer the consequences that arise from it, Internet users today suffer the exploitation of their personal information and the infringement of their privacy rights. Following this analogy, we propose an architecture named "*Fair Personal Information Trades*" (FPIT), that follows the principles of the Fair Trade movement, which offers better trading conditions to, and secures the rights of, marginalised producers and workers [144]. Just like Fair Trade, FPIT can support fair trades of personal information between information owners and information consumers, while protecting their information owners' privacy.

Related Work

One of the first proposed markets for personal information was Laudon's National Information Markets [96], where personal information can be traded through a National Information Exchange. There are two major differences between Laudon's NIM and FPIT. First, in this work we follow a distributed approach, where no third parties are involved for the information exchange. Second, in FPIT the personal information itself is never sold. Temporary access to personal information is sold, by means of appropriate licenses. Information users have to acquire information from their owners each time they need to use it, paying a small fee every time.

A decentralised approach for information markets is Information Crystals [5]. This model aims at creating large groups of personal information, to be used aggregated for data mining, while protecting the owner's privacy.

An interesting work that implements the idea of Personal Information Market (PIM) is [72]. In that approach, only preference and behavioral information is for sale and therefore privacy protection is achieved by keeping Personal Identifiable Information (PII) undisclosed. PII is defined as any piece of information which can potentially be used to uniquely identify, contact, or locate a single person. In our work, both PII and preference/behavioural information can be exchanged. We argue that companies should be able to acquire individuals' contact information, with their consent, for marketing purposes. We address the problem of privacy protection by using data licenses and the convention that no information is allowed to be stored at the company's side. Another important differentiator

is that our platform does not rely on the existence of trusted third parties for transferring personal information.

4.2 Fair Personal Information Trades: Concepts and Architecture

The core principle of FPIT is that the control of personal information should be maintained by its owner. Therefore, companies are not allowed to store individuals' personal information and use it without their consent. The main players in FPIT are the following:

- *Individuals* who voluntarily participate in FPIT, selling access to their personal information.
- *Companies* interested in collecting and processing personal information.

These players are represented in the FPIT architecture by the *entity* component. Both individuals and companies can be called *FPIT-users*, or just *users*.

The resources traded in FPIT are licenses to access personal information of individuals. This renders the task of storing, managing and retrieving personal data a very critical operation in FPIT.

4.2.1 Personal Data Management in FPIT

In order for FPIT to work efficiently, it must contain a privacy enhanced subsystem for the storage of the individuals' personal data. We call this subsystem the "Personal Data Management System" (PDMS). Due to the nature of FPIT, the management of personal data has to meet the following requirements:

- Personal data can only be stored at the owner's side.
- Personal data must always be accessible for licensed use.
- Information security must be ensured and information leaks should be prevented.

A system that satisfies the above requirements for the management of personal information is the Polis platform described in [56]. In Polis, for every individual there is a personal agent, which is constantly accessible over the Internet. The agent contains the personal information, the policies and the contracts of the individual. Each company also has its own agent, which contacts individuals' agents in order to retrieve (some of) their personal information. The functionality

and services of the agents in Polis can be extended by implementing appropriate (cryptographic) protocols. We use this feature in the implementation of the FPIT prototype.

4.2.2 FPIT-Users

Each FPIT-individual entity is characterised by its personal information and its policies. This architecture can be expanded to contain more sophisticated components, like a transaction logging service or a negotiation mechanism.

In order for FPIT to work efficiently, its agents need to have reliable Internet connectivity. This requirement is straightforward for companies. As far as individuals are concerned, constant connectivity is quite common today and is soon expected to become a given. Nevertheless, the protocol for personal information exchange described below could be implemented in such a way, that even if the agent of an information provider loses connectivity, there will be no monetary loss for the companies this user interacted with.

Personal Information Representation

Personal Information traded in FPIT can be Personal Identifiable Information (PII), like the name, phone number, address, birth date etc, as well as preference and behavioral information of a person. In this work we examine the market for trading PII. It is straightforward, though, to expand FPIT to deal with preference and behavioral information as well.

Personal Information can be represented in an XML schema like the one shown in Figure 4.1. This representation is simple and efficient enough to suit the needs of FPIT. Personal information is organised hierarchically in a number of categories, each of which can contain appropriate subcategories. This scheme can be expanded according to the implementation and usage needs of FPIT.

```
<?xml version="1.0" encoding="utf-8" ?>
- <User Description="Personal Data">
  - <Name Description="User's Name">
      <Given Description="Given Name">John</Given>
      <Family Description="Family Name">Doe</Family>
    </Name>
  - <Home-Info Description="User's Home Contact Information">
    - <Postal Description="Home mailing address">
        <Name Description="Name on mailing address">John Doe</Name>
        <Street Description="Home street address">FPIT Street 10</Street>
        <City Description="City">FPIT City</City>
        <StateProv Description="State or Province">FPITia</StateProv>
        <PostalCode Description="Postal Code">11111</PostalCode>
        <Organization Description="Organization Name">FPIT</Organization>
        <Country Description="Country Name">FairInformationTradeLand</Country>
      </Postal>
      <Telecom>...</Telecom>
    </Home-Info>
</User>
```

Figure 4.1: Personal Information in FPIT

Policies and Licenses

Policies are integral components of FPIT trades. Agent policies define whether the agent will accept or reject a transaction request. A policy, represented in an XML schema, contains the following fields:

- *Principals*: The FPIT-entities.
- *Info-item*: Every distinct item of an information provider's personal information.
- *Purposes*: The set of purposes that entitle principals to retrieve data. Some indicative purposes are promotion and statistics. Further additions could be made according to specific transaction needs.
- *Usage restrictions*: Additional restrictions may exist that limit access rights to a specific number of accesses or a specific time interval, or both.
- *Charge*: Value and unit of payment and conditions for charging.

Another important component/concept of this architecture is the license. A license is used to set the rules under which a company is entitled to have access to an individual's personal information. Licenses play a key role in this work, since they are the mechanism that controls personal information use and distribution.

The architecture overview of FPIT is presented in Figure 4.2.

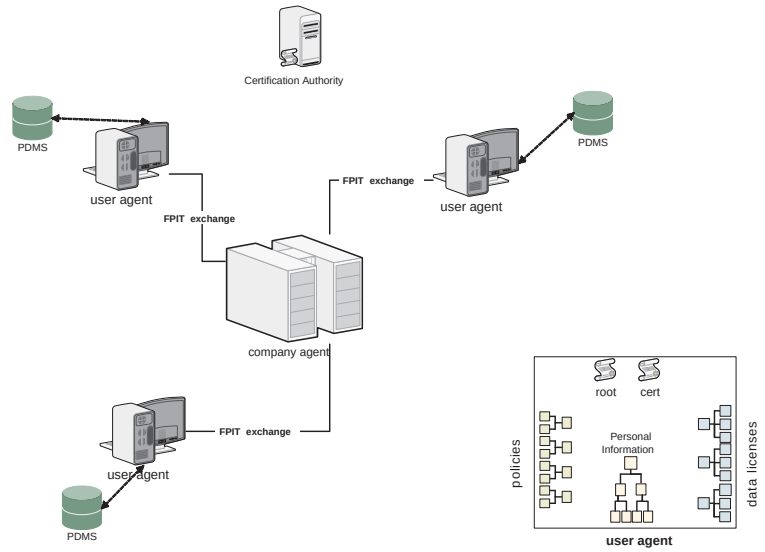

Figure 4.2: The FPIT architecture overview

4.2.3 Payments in FPIT

The payment scheme within FPIT needs to be efficient enough to facilitate large numbers of small amount payments, without entailing substantial transaction costs. Therefore, we consider that micropayments as proposed in [123], suit the aforementioned needs. In this section, we provide a brief description of micropayments, used in FPIT transactions. Micropayments as described by Rivest and Shamir in [123] aim at making successive payments of small amounts of money possible, by eliminating bank/credit card processing costs. Other interesting works on the subject are [115, 104, 41, 10, 147]. As described in Section 4.2.4, micropayments are effortlessly integrated in the Agreement and Purchase protocols.

The main actors in micropayment schemes are Brokers, Vendors and Users. A User becomes authorised to make micropayments by the Broker. A Vendor receives micropayments from authorised users and redeems them through the Broker. Relationships of Users and Vendors with the Broker are long term.

Payword

The micropayment scheme we use in FPIT is Payword, presented in [123]. Payword is a credit-based scheme, based on chains of hash values (Paywords). The Payword protocol is briefly described by the following steps:

1. A User creates an account with the Broker and receives a Payword Certificate. Payword Certificates are used to verify that users are authorised to create valid Payword chains and are renewed periodically (for example monthly).
2. The first time a User interacts with a Vendor, the User computes and signs a *commitment*, containing the root w_0 of a chain of paywords and the Payword Certificate. A chain of paywords w_1, w_2, ..., w_n, is user- and vendor-specific, i.e. it can be used for a series of transactions between the user who created this chain and a specific vendor. Additionally, the user and the vendor must share the same broker in order for a transaction between them to be possible. A payword chain is created by picking the last payword w_n randomly and then computing $w_n = h(w_{i+1})$, for $i = n-1, n-2, ..., 0$. The User sends the commitment to the Vendor.
3. The Vendor receives the certificate and commitment from the User and verifies their signatures. After successful verification, payments can take place, between the User and the Vendor.

4. Each payment from a User to a Vendor consists of a pair (w_i, i), where $i \in \{1, 2, ...\}$, is the number of the i-th transaction. The vendor verifies this payment by using w_{i-1} in a single hash operation.
5. In order to redeem the payments received, the Vendor sends the Broker the last payment received (w_l, l) from each user with the corresponding signed commitment. The Broker charges the User's account the value of l coins and deposits this value into the Vendor's account.

One of the major advantages of Payword is that the Broker does not need to be online in order for a transaction between a User and a Vendor to take place. This efficiency makes Payword suitable for applications, like FPIT, that require incremental payments of small amounts.

4.2.4 Trading Process in FPIT

Locating potential personal information providers can be achieved in several ways. The first and simplest solution is for companies to use their own clientèle database which would contain the agents' contact information of the clients that were interested in participating in personal information trades. Apart from that, other possibilities exist, like the creation of whitepages for participating individuals, or even participating agents' contact information exchanges between companies. In this work, we consider the act of locating the information provider already accomplished and propose a protocol for the actual exchange of the personal information.

As far as pricing is concerned, i.e. the price per personal information item access, this is set to a fixed price of one Payword coin (usually representing the value of one cent). A pricing policy could be used in order to allow individuals to set different prices on their information items. For example, a person's phone number could be more expensive than their age. Prices could also vary depending on the time of day or the season of the year. For example, acquiring one's phone number to call them during the evening or holidays could be more expensive.

Description of Protocols

Once a company (information buyer) finds the contact information of an information provider's agent, the trading process can begin. The information trading process in FPIT consists of two phases: The Initial Agreement phase and the Purchase phase. These phases are described below and depicted in figure 4.3.

During the Initial Agreement phase the following actions occur:

1. The information buyer contacts the information owner, sending a message about the kind(s) of personal information they are interested in, the period

of time for which they are requesting access to the information and the price they are willing to pay for it. For example, an online shop might be interested in a person's e-mail for one year in order to send them promotional e-mails with offers and be willing to pay one coin for each e-mail.

2. The information owner's agent receives the request, checks whether it complies with its policies and responds accordingly.

3. If the request is accepted, the information buyer agent sends the commitment M, according to the Payword protocol (Section 4.2.3).

4. The information owner verifies the buyer's certificate according to the Payword protocol.

5. If the verification is successful, a license is sent to the buyer, entitling them to the requested access to the owner's personal information.

After having established the initial agreement with the data owner, the data buyer can make several purchases, according to the agreed upon license. During a Purchase phase, the following actions take place:

1. The data buyer requests a specific item of personal information.

2. The owner's agent receives the request and verifies the accompanying license.

3. If the license is valid, an ACCEPT message is sent to the buyer's agent (verifying at the same time that the owner's agent is up).

4. The buyer sends the payment for the requested items according to the Payword protocol.

5. The owner's agent sends the requested information.

Using this protocol companies are protected from potential malicious information providers. The access to the personal information is not prepaid at the Initial Agreement phase and thus, information providers cannot receive their payment and disappear. Payment occurs each time an information item is requested. Therefore, the company can confirm that the information provider's agent is up before making any payments. The only way for an information provider to cheat is to receive the coin for the particular information item requested and then disappear. Even then, the gain for the information provider as well as the loss for the company will be minuscule. Besides, the company can always revoke the stolen Payword coin at the Broker.

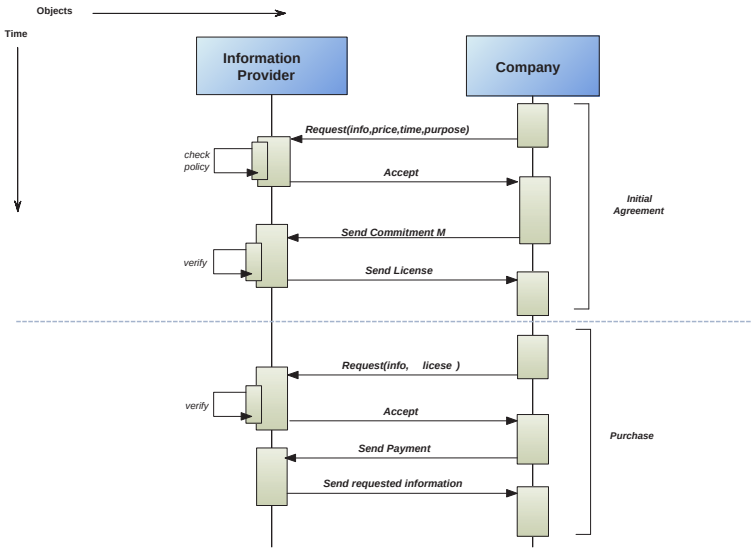

Figure 4.3: FPIT sequence diagram for transaction protocols

4.3 The FPIT Prototype

We implemented an FPIT prototype and performed proof-of-concept experiments. The main objective was to become acquainted with the practical difficulties intrinsic to a platform like FPIT. The prototype is implemented in *Java*. We use the *Bouncycastle* library for the cryptographic primitives. The management of the personal data is done with the Polis platform developed in [56]. The Payword micropayments used in the prototype are an adaptation of the Payword implementation in [64].

The development of the prototype proved to be straightforward. We used the following components for the experiment: An FPIT-shop, two FPIT-individuals and a Broker. The scenario of the experiment was that a shop named *shopfpit* buys and retrieves personal information from two FPIT-individuals, *alicefpit* and *bobfpit*. A snapshot of an agent used in FPIT is given in Figure 4.4.

First, a pairing is established between the shop and each of the individuals. The outcome of this pairing is a digital contract between the shop and the individual. The contract defines which information items can be retrieved, under which terms, and at what cost.

Then, each time the shop wants to retrieve personal information from an affiliated individual, the FPIT-shop must create Payword coins or use already existing coins, and send the coins and its request to the individual's agent. The agent checks the coins, checks the request and provides the requested data.

In the FPIT prototype, we employ some security measures: the communication

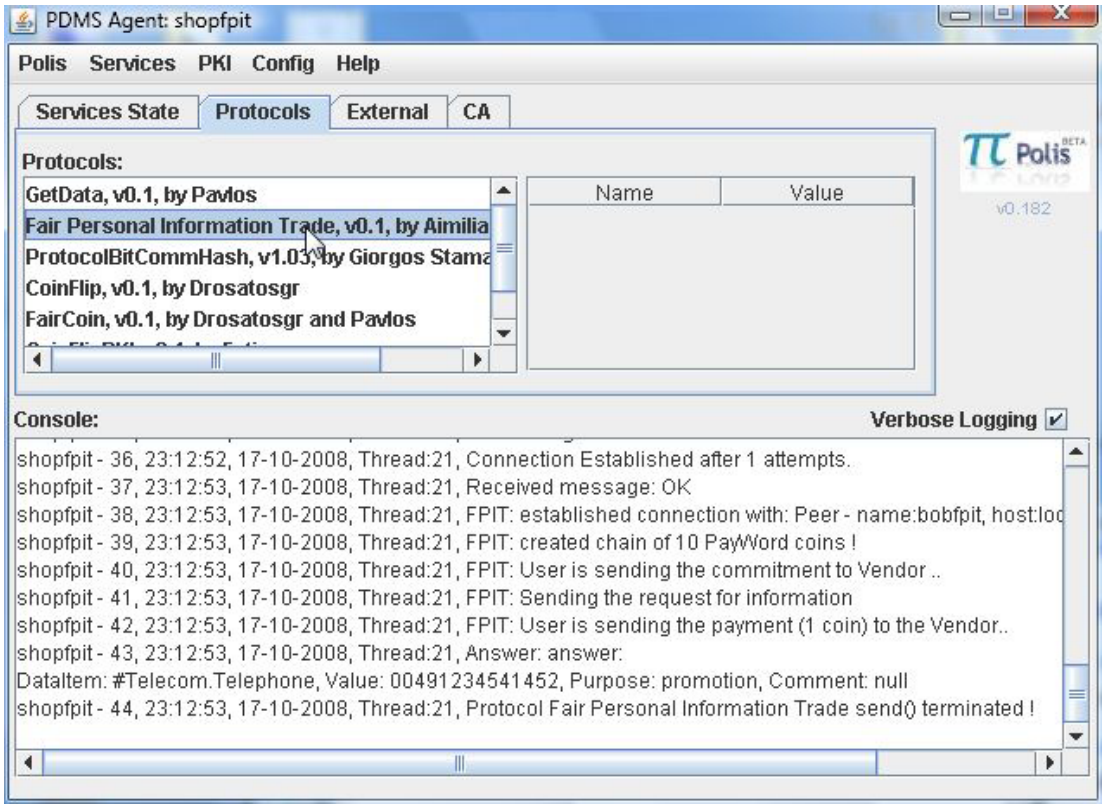

Figure 4.4: FPIT agent snapshot

is performed over SSL sockets with both server and client authentication enabled. Contracts are digitally signed by both parties. However, at this stage we did not address malicious behavior or general fault tolerance issues. We are mainly interested in verifying the basic operations involved in the FPIT transactions and their efficiency. The main outcome of the experiments is that all the building blocks of the FPIT-platform work well. This confirms that the basic functionality of FPIT is attainable.

4.4 Discussion

We believe that Fair Personal Information Trades provides a straightforward proposal that could constitute a regulatory solution to the large scale privacy invasion that is currently perpetrated. More specifically:

- FPIT sets well-defined, clear rules for the proper utilization of personal information, giving the ability to individuals to control and decide upon how information about themselves is used.

- It provides a legitimate and fair, yet efficient enough, way for companies to acquire dynamic, up-to-date information, relevant to the purpose of their intended use.
- In general, FPIT puts the idea of compensating the information owner into practice. It ensures that individuals who provide information are fairly compensated for their service, while their personal information is being protected, at least as much as it is protected currently, but most possibly even more.
- It attempts to reduce information collecting companies' unrestrained personal information exploitation and motivates them to behave more responsibly.
- It successfully combines the use of data licenses with the ideas of information markets and micropayments to propose a broader solution for trading personal information.

Chapter 5

Television Audience Measurements: A New Privacy-Preserving Approach Utilizing Smart TVs

5.1 Introduction

Internet-enabled television systems (Smart TVs) are a recent development in television and home entertainment technologies. In this work, we propose a privacy-preserving approach for computing Television Audience Measurement ratings, utilizing the capabilities of Smart TV technologies. Smart TVs communicate over the Internet to calculate aggregate measurements, without the need for proprietary hardware. Contemporary cryptographic primitives are utilized to ensure the privacy of the individuals and the ratings' validity. User compensation capabilities are introduced to bring some of the data processing company profits back to the data owners. Experimental results of the Android-based prototype implementation illustrate the feasibility of the approach.

The internetworked environment nowadays extends beyond computers, smartphones and tablets, including a multitude of electronic devices. *Smart TVs*[1] are the next logical step in television technology and the term is used to describe the integration of the Internet and Web 2.0 features into modern television sets and set-top boxes, as well as the technological convergence between computers and these television sets/set-top boxes. As the technology improves, many of these sets are becoming as capable as standard computers when it comes to web browsing and even Internet video.

[1] http://en.wikipedia.org/wiki/Smart_TV

This combination of traditional TV functionality with computational and networking capabilities, makes Smart TV technology capable of a whole new set of applications. This constant network accessibility and elevated ability for data collection raises increased security and privacy concerns amongst users, who can now be monitored by multiple devices and tracked by a multitude of data items. The ability to combine the data collected from these sources increases the users' concerns and enhances the need for advanced trusted computing technologies [150, 78].

Television is still nowadays one of the dominant mediums for information and entertainment. Information about television audiences provide valuable insights to broadcasters and the advertising industry on recent trends. Television Audience Measurement (TAM) systems aim at calculating qualitative and quantitative TV audience measurements.

The predominant solution for television audience measurements today is the Nielsen People Meter[1]. A small representative viewer sample is selected to participate in the measurements and specialised metering devices, are installed at the participants' side. The measured viewership data from the metering devices is transferred every night directly to the company's servers. This data is processed and, using statistical inference, extrapolated to reflect the expected viewership percentage for the general population. Larger participant groups are unadvisable in this measuring scheme, due to the equipment and the participation compensation costs entailed.

The reliability of the People Meter measurements is often questioned, both due to the possibility of an ill-selected viewer sample and statistical inference mechanism, as well as the potential tampering of the results by the measuring company. The company is often accused by television networks, minority groups and even legal entities of misrepresenting minorities in their ratings.

This measurement process raises important privacy issues for the participants. A person's viewing record can reveal sensitive information about the person's preferences and habits. A privacy-preserving method for creating accountable TAMs is needed, in order to utilize television ratings information, while protecting the participants' privacy.

Apart from the trust, necessary to achieve users' participation, economic aspects of the IoT applications arise. User information processing can bring about significant profits to companies, therefore the users should be able to control the amount of the disclosed information and its uses, as well as be able to receive some compensation for the service they offer. In the field of Television Audience Measurements in particular, the processing of the viewership data brings about important financial effects, e.g. determining the advertisement pricing depending

[1] http://en.wikipedia.org/wiki/People_meter

on the popularity of the show and the earnings of the show participants.

In this work, we present PrivTAM, a system for calculating privacy-preserving TAMs through Smart TV technology. The core of PrivTAM is a privacy-preserving cryptographic protocol, which accepts as input the viewing records from users' Smart TVs and performs secure multi-party computations [148, 21] to calculate the TAMs. Additionally, functionality for the compensation of the participants is supported.

We consider privacy-preserving television audience measurement a similar problem to e-voting protocols, since both protocols have similar requirements to achieve accountability and privacy goals. We present the main e-voting protocol requirements [113, 18, 87, 94, 42] and explain how these requirements are applied in the PrivTAM case:

- Privacy - all individual viewership records must be secret. This is an important premise of this work. No entity, including the measurement service, should be able to determine the viewership records provided by an individual.
- Completeness - all valid records must be counted correctly.
- Soundness - dishonest records cannot disrupt the measurement process.
- Unreusability - no user can submit their record more than once.
- Eligibility - only those who are allowed to participate can submit their records.
- Verifiability - nobody can falsify the result of the TAM process and it is possible to verify, if necessary, that the measurement result has not been tampered.

In Sections 5.4.4 and 5.5.4 the ways the above requirements are satisfied in the PrivTAM protocol are discussed. Additional requirements of e-voting systems include "verifiably cast-as-intended", which means each participant should be able to verify their record has been accurately counted in the results. This feature is not supported in PrivTAM, since the encrypted individual viewership records are not published on public bulleting boards, as in e-voting protocols. Although such a feature would be possible to be added to the PrivTAM functionality, we consider this unnecessary (since audience measurement results are not as crucial as election results) and is omitted for simplicity.

The computations of PrivTAM are performed between software agents, which are located at the participants' Smart TVs, and a Trusted Authority (TA). Each Smart TV contains an agent which continuously collects the viewing records of

its owners. The Trusted Authority coordinates the TAM computation, verifies the validity of the records, collects the encrypted results and provides the compensation to the participants. This process is performed using encrypted viewing records, hence the record contents are never revealed to the Trusted Authority.

In addition to the main functionality, we created an extension to the PrivTAM protocol, where the demographic information submitted by the participants is certified and therefore can be checked not only for the validity of its form but also for the accountability of its content. Finally, we develop a prototype implementation and perform experiments that confirm the feasibility of the approach.

Some of the advantages of our approach in comparison to traditional TAM systems are:

- The privacy of participants' viewing records and demographic information is preserved.
- The practically unrestricted number of participants that produce the PrivTAM results. can lead to more reliable TAM measurements.
- Fine grained measurements are supported, which can be periodically calculated in small time intervals as well as specific one-time queries.
- The cost for conducting a TAM is reduced, since no specialised equipment is required and only the viewers participating in a calculation are (optionally) compensated.
- Measurements using records from any Internet-enabled broadcasting medium (e.g. Broadcast TV, Cable TV, IPTV and Satellite TV) are supported.

Our solution requires Smart TVs to have permanent Internet access, a requirement which is considered reasonable nowadays. Moreover, the computational and networking requirements of PrivTAM can be easily fulfilled by modern embedded Smart TV platforms.

5.2 Related Work

With the recent growth of IPTV and on demand media services, new audience measurement schemes were proposed, mainly to support multicast media delivery, contrary to the traditional television broadcast model [8, 149]. These audience measurement schemes gather the viewership data either at the viewer side, using video on demand set-top boxes, or at the network side, using the network media delivery infrastructure (Ethernet switches / DSLAMs).

These works bear similarities to the PrivTAM approach in that they are scalable (i.e. a larger amount of viewers can participate compared to the People

Meter) and measurements can be conducted periodically or on demand, with measurement intervals defined dynamically. Additionally, in PrivTAM, as in network based measurement schemes [149] no special equipment, apart from the SmartTV is required for the measurement to take place. In set-top box based measurement schemes [8], the media provider's set-top boxes need to be utilized and maintained for the measurement process. The PrivTAM scheme is mainly concerned with real time, television broadcast measurements on Smart TVs, although on demand media can also be supported.

To our knowledge this work is the first to propose a privacy-preserving scheme for conducting scalable and accountable audience measurements. Previous works identify the privacy issues of audience measurement schemes, but do not take particular measures from protecting the participants' privacy during the measurement process. Additionally, the measurement results in PrivTAM can be validated, making it difficult for the measuring service to tamper with the TAM results. Moreover, the enhanced version of the protocol (Section 5.5) supports functionalities beyond plain TAM measurements, containing additional verified information about the participants (e.g. health or opinion data), while still preserving their privacy.

The sensitivity of the viewing records is stressed by both the Video Privacy Protection Act [136] and the Cable TV Privacy Act [135]. In general, TAMs are products of aggregation operations and therefore our work is related to common privacy-preserving aggregation systems. Overall, we consider that PrivTAM lies between privacy-preserving aggregation systems and e-voting systems, offering verifiable, privacy-preserving, aggregation functionalities. Privacy-preserving data aggregation in people-centric urban sensing systems is discussed in [128].

The PrivTAM computation, uses encrypted viewership vectors that contain, apart from the viewership record, demographic information about the viewer (e.g. age group, gender) to facilitate viewership statistics. The main PrivTAM functionality is enhanced by incorporating modern cryptographic technologies to ensure the accountability of encrypted demographic information contained in the viewership vector. The use of cryptographic primitives, including secure multi-party computations in particular, to create privacy-preserving applications is also analyzed in [61]. Additional cryptographic technologies used in this work include anonymous credentials [31, 27], blind signatures [37, 30] and zero-knowledge protocols [69, 68].

Moreover, PrivTAM takes into account the economic aspects of privacy[139, 2, 88, 92] and supports compensation Electronic cash services and micropayments schemes [123, 104] facilitate user compensation functionality in online environments. Micropayment schemes have been proposed many years ago, but they have not had the expected success and practical applicability to ensure widespread adoption. Experts have expressed doubts as for the chances of success of micro-

payments, due to lack of efficient implementation mechanisms as well as lack of market adoption incentives [100, 110]. However, virtual currencies, such as BitCoin [107], an anonymous online payment system and electronic currency, are becoming increasingly used and accepted today, while micropayment technologies are beginning to come back into action [145] today with the new technological advancements and the spread of the IoT as analyzed in [141, 138].

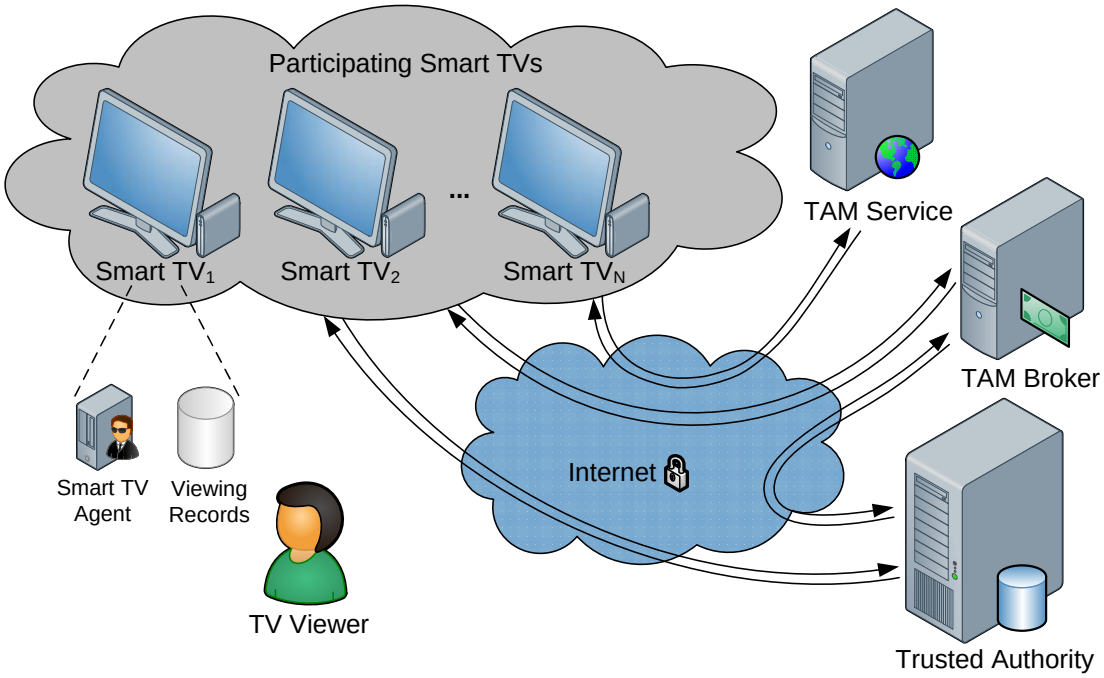

Figure 5.1: The PrivTAM architecture.

5.3 The PrivTAM System

An overview of the PrivTAM system architecture, built on top of Smart TV technology, is shown in Figure 5.1. The main parts of the architecture are the participating Smart TVs, the Television Audience Measurement Service (TAM Service) and the Trusted Authority (TA). Every Smart TV contains a software agent that collects and stores the Smart TV's viewing records and maintains a set of demographic data, such as gender, age and educational level of the viewers. The agent manages the viewers' personal data, provides controlled access to the data, and has the ability to participate in distributed protocols and computations.

The TAM Service collects the measurements and is responsible for coordinating the distributed key generation [108] for the public-key cryptosystem between

itself and a group of L TV agents. These L agents are chosen with a verifiable random selection [54] and participate in both the public-key creation phase and the result decryption phase. The selection of the set of L agents can be repeated on a regular basis, for example every one or few days. With each execution a new set of agents will have the responsibility of the distributed key generation process. The verifiability of the agents' random selection can be found in Section 5.4.3.

The TA is responsible for coordinating the PrivTAM computation process. Time is divided into consecutive intervals, and for each interval, an aggregate result is periodically calculated using input from the participating Smart TVs. Many conventional TAM systems, report audience results in intervals of 15 minutes. PrivTAM can obtain equivalent results by using the same interval value. The TAM Broker is used to facilitate the (optional) compensation functionality described in Section 5.4.3.

Each Smart TV agent encrypts its viewership vector with the public-key of the measurement and sends it to the TA. The validity of the the encrypted vector is verified using the cryptographic protocol described in Phase 2 of Section 5.4.3. Following a successful verification, the TA adds this vector to the current encrypted result of the measurement. The final encrypted TAM rating is transmitted to the participants in the distributed key generation for decryption and next the result is revealed to the TAM service.

Apart from the main TAM computation, performed in consecutive intervals, the PrivTAM system can support more specific requests for TAM ratings, since the TV agents have the full viewing records of their owners. Some examples of such specific TAM ratings are:

- How the TAM of a channel changed at a specific point in time, e.g. when some person (actor or politician) was screened in a TV Show.
- What channels were watched by the audiences of a particular channel after a specific time.
- What was the health state of the viewers who watched a show. This example requires some extra personal health information, apart from the demographic data, that is possible to be kept in the TV agent (an example of privacy-preserving healthcare statistics can be found in [53]).

5.4 The PrivTAM Protocol

In this section, we present the main cryptographic protocol used in PrivTAM. Communication between the entities in our protocol is performed over secure sockets (SSL/TLS) to assure the confidentiality and the integrity of the communication, but also server authentication and client authentication whenever

applicable. During the computation the individual user's information is not disclosed in any stage of the process, only the aggregate results are revealed at the end.

5.4.1 Security Model and Privacy

We first define the security model and the degree of privacy achieved and then proceed with the problem definition and the outline of the computation. We show that the proposed protocol is safe in the Malicious Model, assuming that the TA is honest but curious (HBC) and the TAM Service is the final receiver of the TAM results (see Section 5.4.4).

Regarding privacy, there are two distinct problems that arise in the setting of privacy-preserving computations [101]:

- The first is to decide which functions can be safely computed, where safety means that the privacy of the participants is preserved if the result of the computation is disclosed. We will assume that the computation outcomes do not violate the privacy of the participating viewers and will not further consider this problem in this work.
- The second is how to compute the results, i.e. which algorithms and protocols to use, to minimize the damage to privacy. For example, it is possible to pool all the unencrypted individual viewership vectors in one location and run the computation algorithm on the pooled data. However, this would compromise the privacy of the viewership vector owners.

In this work we focus on the latter problem, attempting to address the question of how to compute the TAM ratings without gathering the viewing records in one location, and in a way that reveals nothing but the aggregate TAM result of the computation.

5.4.2 Problem Definition and Privacy Goals

We define the PrivTAM problem for validated, privacy-preserving TAM ratings. A PrivTAM problem instance consists of:

- **N Smart TVs** - $TV_1, TV_2, \ldots, TV_N$ and the viewing records of their owners.
- **The PrivTAM computation:** The software agents of the Smart TVs participate in a privacy-preserving computation.
  - **Input:** The viewership vector of each participant.

– **Output:** The aggregate TAM measurement for the participating viewership vectors.

We assume that one viewership vector is submitted per Smart TV. We do not consider user identification issues amongst Smart TV users, as this is an existing issue in TAM systems and is out of the scope of this work.

The privacy goal of PrivTAM is that no information about the viewership record and the demographic data of any individual viewer should be disclosed, beyond what can be inferred from the aggregate result or what is already known.

The main entities that participate in a PrivTAM computation are the viewers (which submit their viewership data) through their Smart TVs, the trusted authority (TA) and the TAM service. The critical data of the computation are the viewership records and the demographic data of the viewers. In the security analysis of the protocol we will show that it is fairly robust: The privacy goal is assured, on the condition that the TAM service and the trusted authority are (at least) honest but curious, whereas there is no particular assumption on the viewers; the protocol can handle even a number of malicious viewers, who may try to manipulate the aggregate result.

5.4.3 Outline of the Computation

The computation consists of three main phases. In Figure 5.2, the participating entities of each phase are illustrated. The full descriptions of the three phases are given in the following paragraphs.

- In Phase 1 a distributed key generation for a Threshold Paillier Cryptosystem is performed.
- In Phase 2 the privacy-preserving TAM computation takes place.
- In Phase 3 the final encrypted TAM is forwarded for decryption and the result is revealed.

Phase 1. During Phase 1 the TAM Service selects an L-sized subset of the N-sized set of all the participating TV agents with a verifiably random process. An example of a publicly verifiable random selection process is described in [54]. As a result, L agents are randomly selected and the TAMS can prove that they are really random agents. This technique prevents the TAM Service from making a biased or impeachable group selection. Then, the TAM Service and the L selected TV agents execute a cryptographic protocol for the distributed key generation of the Threshold Paillier Cryptosystem [43]. We use the following Threshold Decryption Model, which is an adaptation of the corresponding definition in [18] to our

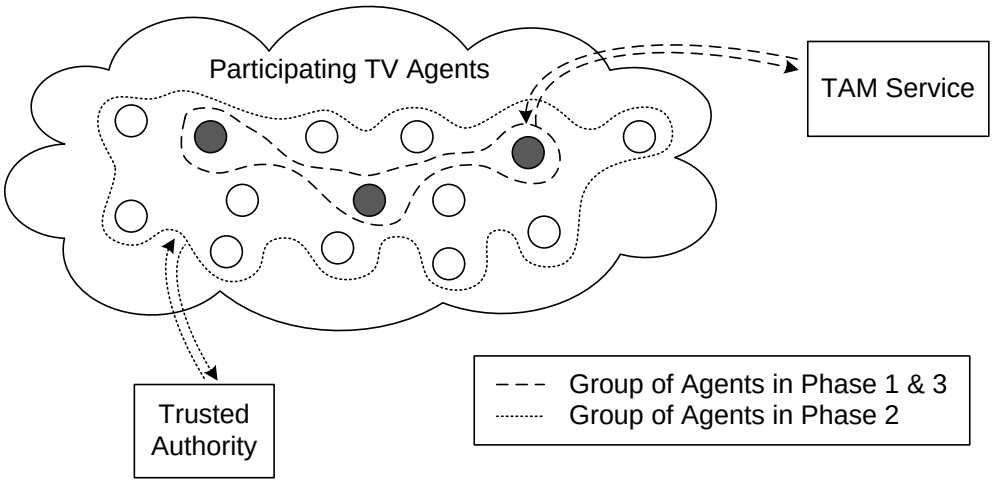

Figure 5.2: Illustration of protocol participants.

needs, so that the distributed key generation can be performed without a trusted dealer [108]. Avoiding the need for the trusted dealer makes the PrivTAM system more relevant for the TAM application context.

Definition 1 (Paillier Cryptosystem) *The Paillier cryptosystem is a probabilistic asymmetric cryptographic algorithm for public key cryptography. Its security is implied by the Decisional Composite Residuosity Assumption (DCRA) [111].*

Definition 2 (Threshold Decryption Model) *In a threshold cryptosystem, instead of merely decrypting the encrypted message, n parties P_i are used with their secret keys, so that at least t parties, where $t \leq n$, are required to decrypt the message. The decryption process includes the following players: a combiner (can be one of the n parties), a set of n parties P_i, and users. We consider the following scenario:*

- *During the initialisation phase, the parties use a distributed key generation algorithm to create the public key PK of their private keys SK_i. Next the parties publish their verification keys VK_i.*
- *To encrypt a message, any user can run the encryption algorithm using the public key PK.*
- *To decrypt a ciphertext c, c is forwarded to the combiner and n parties. Using their secret keys SK_i and their verification keys VK_i, each party runs the decryption algorithm and outputs a partial decryption c_i with a proof of validity of the partial decryption $proof_i$. Finally, the combiner uses the combining algorithm to recover the cleartext, provided that at least t partial decryptions are valid.*

In PrivTAM, we use the Paillier public key generated in Phase 1 for the encryption of the viewership vectors and utilize the Pailler Cryptosystem's homomorphic property in Phase 2. In addition, we specify that t is equal to n in our Threshold Decryption Model, meaning that all the parties are required to decrypt a message. Setting $t = n$ is important to ensure that the final result cannot be decrypted without the active participation of the TAM Service. Phase 1 should be repeated occasionally, to renew the keys and the set of L agents.

Phase 2. During this phase, the TA coordinates the measurement process, collects and validates the encrypted viewership vectors of the participants. Upon successful validation, the TA adds the submitted viewership vector to the current TAM result, and sends the compensation to the participant.

In detail, Phase 2 begins with the TV agents that hold viewing records for the particular time period, creating their viewership vectors (Figure 5.3). Each such vector is submitted to the TA for validation. The validation process is based on a zero-knowledge proof (ZKP) that an encrypted message lies in a given set of messages [18]. This way, when encrypting a message, it is possible to append a proof that the message lies in a public set $S = \{m_1, \cdots, m_q\}$ of q messages without revealing any further information. This proof is described in detail in Section 5.4.4.

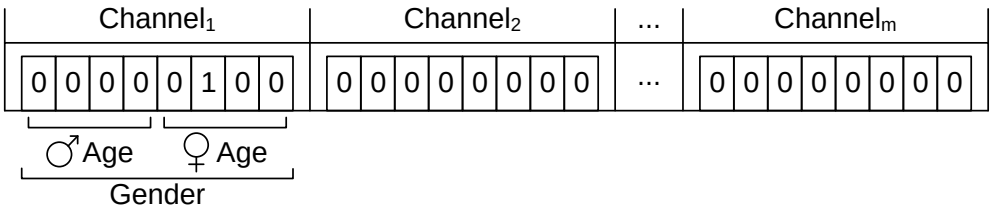

Figure 5.3: Example of a viewership vector.

In Figure 5.3, the viewership vector for m TV channels is illustrated. The vector section for each channel consists of a number of ciphertexts, which result from the number of demographic data items used in the vector. Every demographic data item is respectively represented by some ciphertexts whose number depends on the number of parts that the demographic element is divided. Additionally, in our case the two demographic data items are used as dependent parameters, in order to be able to draw conclusions about combined characteristics of the viewers and are hence represented as one combined vector. Of course, it is possible to use independent parameters for the demographic data items in our measurements.

In our example, the represented data items are the age group and the gender of the viewer. The gender categories are male and female and the age groups are "$Age_1 \leq 24$", "$25 \leq Age_2 \leq 40$", "$41 \leq Age_3 \leq 55$" and "$Age_4 > 55$". Consequently, a combination of 8 ciphertexts is created to represent these data items.

In order to indicate the channel the viewer was watching, the representation of their demographic data is added to the viewership vector section for the corresponding channel. For example, in our case the viewer's gender is female, her age is between 25 and 40 years old and she was watching $Channel_1$. All the values in the viewership vector lie in the public set $S = \{0, 1\}$ and they are encrypted using the public key that is generated in Phase 1. Every participant should prove that their vector is valid so that the TA can avoid any malicious behavior from them. More specifically, the participants should prove that:

1. Every ciphertext in the viewership vector should lie in the set $S = \{0, 1\}$.

2. The multiplication of the ciphertexts in every channel should lie in the set $S = \{0, 1\}$.

3. Finally, the multiplication of all ciphertexts in the viewership vector should be equal to "1". This means that the participant was watching TV (i.e. is not providing a null viewership vector).

The multiplication of the ciphertexts in the above proofs utilizes the homomorphic property of the Paillier Cryptosystem [111]. In Table 5.1, the details of the levels of the above proofs are illustrated. In our case, where dependent demographic data items are used, the 2^{nd} level of proofs is redundant. However, it remains in the protocol to cover the cases that the demographic data items are used as independent parameters.

Table 5.1: The levels of proofs for a viewership vector.

	$Channel_1$	$Channel_2$	...	$Channel_m$
Ciphertexts of viewership vector	$E_{1,1}$ $E_{1,2}$ ... $E_{1,\delta}$	$E_{2,1}$ $E_{2,2}$ ... $E_{2,\delta}$	...	$E_{m,1}$ $E_{m,2}$ ... $E_{m,\delta}$
1^{st} Level Proofs	$P_{1,1}$ $P_{1,2}$ ... $P_{1,\delta}$	$P_{2,1}$ $P_{2,2}$ ... $P_{2,\delta}$		$P_{m,1}$ $P_{m,2}$ ... $P_{m,\delta}$
Multiplication of ciphertexts per channel	$E_1 = \prod_{d=1}^{\delta} E_{1,d}$	$E_2 = \prod_{d=1}^{\delta} E_{2,d}$	...	$E_m = \prod_{d=1}^{\delta} E_{m,d}$
2^{nd} Level Proofs	P_1	P_2		P_m
Multiplication of all ciphertexts	$E = \prod_{ch=1}^{m} \prod_{d=1}^{\delta} E_{ch,d}$			
3^{rd} Level Proof	P			

1^{st} and 2^{nd} Level of ZKPs: Ciphertexts should lie in the set $S = \{0, 1\}$.
3^{rd} Level of ZKP: Ciphertext should equal to "1".

Definition 3 (Homomorphic Encryption) *Homomorphic encryption [122, 63] is a form of encryption where one can perform a specific algebraic operation on the plaintext by performing a (possibly different) algebraic operation on the ciphertext.*

The additive homomorphic encryption property of the Paillier cryptosystem means that the multiplication of the encrypted values corresponds to the addition of the decrypted ones. Let x_1 and x_2 be two plain integers, $x_1, x_2 \in \mathbb{Z}_{n_p}$, and (n_p, g) the Paillier public key. If r_1 and r_2 are two random numbers such that $r_1, r_2 \in \mathbb{Z}^*_{n_p}$, then the encryption of message m is $\mathcal{E}(m) = g^m r^{n_p} \bmod {n_p}^2$, and the Paillier homomorphic property holds, since:

$$\begin{aligned} \mathcal{E}(x_1) \cdot \mathcal{E}(x_2) &= (g^{x_1} \cdot r_1^{n_p}) \cdot (g^{x_2} \cdot r_2^{n_p}) \\ &= g^{[x_1 + x_2 \bmod n_p]} \cdot (r_1 r_2)^{n_p} \bmod {n_p}^2 \\ &= \mathcal{E}([x_1 + x_2 \bmod n_p]) \end{aligned} \tag{5.1}$$

Once the viewership vector is validated by the TA, the vector is multiplied, using the homomorphic property, with the current TAM result. More specifically, every ciphertext of the viewership vector is multiplied with the corresponding ciphertext of the current TAM total by taking advantage of the homomorphic property. We assume that the TA only logs the participants in a measurement in order to ensure unreusability of the vectors. However, even if the vectors were stored, the TA would not be able to reveal their contents, unless all the participants of the threshold decryption are malicious and collude towards this purpose. The final result of Phase 2 is the encrypted aggregate TAM of the particular query, which ensures k-anonymity (see Section 5.4.4), where $k = N$ and N is the number of all the participants who take part in the TAM.

Payments in PrivTAM. The PrivTAM system can support functionalities for the compensation of participants, either in the form of financial payments or in the form of vouchers or credit. The requirement for a participant to be compensated is that they provide a valid viewership vector to the computation. After the successful verification of a participant's viewership vector, the TA sends the compensation to the participant.

The compensation functionality introduced in the PrivTAM system, aims at addressing the negative externalities that arise from the exploitation of the viewers' information. Audience measurements have significant economic impact on the market and bring about important profits to the companies involved (i.e. measurement, broadcasting, advertising companies). These companies profit from the use of the viewers' data and we believe it is appropriate that they are compensated for their contribution and that the data remain under the owners' control. The compensation functionality offers the possibility of a fairer data exchange

between viewers and metering companies and incentivises users to participate. Additionally, the privacy-preserving feature of the compensation mechanism, reassures participants of the protection of their data and allows them to maintain control over them. The shared data cannot be linked back to them or misused in the future.

However, the introduction of the compensation functionality into the PrivTAM system can add some computational overhead, especially in the case of financial compensation. In that case, the payment scheme within PrivTAM needs to be efficient enough to facilitate large numbers of small amount payments, without entailing substantial transaction costs. Therefore, we draw techniques along the lines of micropayments, as proposed in [123].

The main actors in micropayment schemes are Brokers, Vendors and Users. A User becomes authorised to make micropayments by the Broker. A Vendor receives micropayments from authorised users and redeems them through the Broker. Relationships of Users and Vendors with the Broker are long term. In PrivTAM, the Smart TV owners can act as Vendors and the TA can act as a user making micropayments. The TAM Broker is introduced in the architecture to facilitate the payments (Figure 5.1).

An example of a micropayment scheme suitable for PrivTAM is Payword, presented in [123]. Payword is a credit-based scheme, based on chains of hash values (called Paywords) and the Broker does not need to be online in order for a transaction between a User and a Vendor to take place. This fact makes Payword suitable for applications, like PrivTAM, that require incremental payments of small amount payments. Other micropayment schemes can be supported, following the same principles of operation.

Alternatively, non-monetary compensation, including credit that can be redeemed with participating companies, or even services within the SmartTV environment, can be offered to participants. Another possible disadvantage of the compensation functionality is the cost increase for measurement companies, as each participant receives at least a small compensation for their participation. However, this cost could be absorbed by the companies who benefit from audience measurements, e.g. broadcasting and advertising companies. Apart from directly covering the financial cost of the participants' compensation, these companies could also supply credit or services for the measurement company to give to the participants. The amount of compensation for each PrivTAM calculation is fixed for simplicity, but methods for providing different pricing could be introduced into the system. It is important to stress that the collected points of each participant are not recorded in a profile by a centralised service, but are kept at the participant's side.

Phase 3. In Phase 3, the final encrypted result of Phase 2 is forwarded to the

L selected TV agents of Phase 1 and the TAM Service. The L agents perform partial decryptions and send the results to the TAM Service which acts as the final participant and combiner of the threshold decryption. This way, only the TAM Service can see the final result of the calculation, which is acceptable if the TAM Service is considered honest and reports accurately the decrypted result. In order for the PrivTAM calculation to be protected from inaccurate reporting of the results from the TAM Service, a verification mechanism can be introduced to validate the announced results. This verification could be accomplished by using multiple combiners in the threshold decryption, to confirm the announced results from the TAM Service.

5.4.4 Security Discussion

In this section, we will show that the proposed protocol fulfills the requirements of the basic PrivTAM problem described in the introduction, i.e. privacy, completeness, soundness, unreusability, eligibility, and verifiability. The security model holds for Malicious viewers, with the assumption that the TA and the TAM Service are honest but curious, that these entities do not collude, and that the TAM Service is the final recipient of TAM results. Malicious viewers can submit any value as input to the computation or even abandon the protocol at any step. A definition of the Malicious Model is given in [91] and a more detailed description in [67]. An honest but curious party (adversary) [3] follows the prescribed protocol properly, but may keep intermediate computation results, e.g. messages exchanged, and try to deduce additional information from them other than the protocol result.

Some basic facts about the protocol are: The viewership vectors submitted by the TV agents is encrypted with the Paillier cryptosystem [111], which offers Semantic Security [70], that is, it is infeasible for a computationally bounded adversary to derive significant information about a message (plaintext) when given only its ciphertext and the corresponding public encryption key. More details about the semantic security of the Paillier scheme (known and as Decisional Composite Residuosity Assumption (DCRA)) can be found in [111].

Another fact is that the final TAM result is obtained from the encrypted aggregate result using a Threshold Pailler cryptosystem. A security analysis of the Threshold Version of the Paillier Cryptosystem is given in [108]. The Threshold Paillier cryptosystem and its homomorphic property ensure that the viewing records are never disclosed and cannot be associated with any particular participant.

Finally, the validity of each viewership vector is proven by a zero-knowledge proof, which does not leak any information about the actual value of the viewership data beyond the fact that the vector is valid. In particular, the following

zero-knowledge proof is used in the validation process in Phase 2. The security analysis of this zero-knowledge proof is given in [18].

Zero-Knowledge Proof that an encrypted message lies in a given set of messages [18]. Let n_p be a ℓ-bit RSA modulus, $\mathcal{S} = \{m_1, \ldots, m_q\}$ a public set of q messages, and $c = g^{m_i} r^{n_p} \mod n_p^2$ an encryption of m_i where i is secret. In the protocol, the prover P convinces the verifier V that c encrypts a message in $\mathcal{S}$.

1. P picks at random ρ in $\mathbb{Z}_{n_p}^*$. He randomly picks $q-1$ values $\{e_j\}_{j \neq i}$ in $\mathbb{Z}_{n_p}$ and $q-1$ values $\{v_j\}_{j \neq i}$ in $\mathbb{Z}_{n_p}^*$. Then, he computes $u_i = \rho^{n_p} \mod n_p^2$ and $\{u_j = v_j{}^{n_p}(g^{m_j}/c)^{e_j} \mod n_p^2\}_{j \neq i}$. Finally, he sends $\{u_j\}_{j \in \{1,\ldots,q\}}$ to V.

2. V chooses a random challenge e in $[0, A[$ and sends it to P.

3. P computes $e_i = e - \sum_{j \neq i} e_j \mod n_p$ and $v_i = \rho r^{e_i} g^{(e - \sum_{j \neq i} e_j) \div n_p} \mod n_p$ and sends $\{v_j, e_j\}_{j \in \{1,\ldots,q\}}$ to V.

4. V checks that $e = \sum_j e_j \mod n_p$ and that $v_j^{n_p} = u_j (c/g^{m_j})^{e_j} \mod n_p^2$ for each $j \in \{1, \ldots, q\}$.

We note that r is the random number which was used for the encryption of message m_i and $a \div b$ is the quotient in the division of a by b. According to Theorem 2 of [18], it holds that T iterations of the above protocol is a perfect zero-knowledge proof (against an honest verifier) that the decryption of c is a member of $\mathcal{S}$, for any non-zero parameters A and T such that $1/A^T$ is negligible.

Based on the above facts we can show that the protocol preserves the privacy of each viewership vector by satisfying the privacy criterion of k-anonymity.

Definition 4 (k-anonymity) *An simple definition of k-anonymity [38] in the context of this work is that no less than k individual users can be associated with a particular personal viewing record.*

In other words, in PrivTAM, unless some background information on the viewership information is available, an adversary cannot associate less than N viewers with any particular vector. This means that N-anonymity is assured for the viewers.

A summary of the critical data items and the main entities of the basic PrivTAM protocol is presented in Table 5.2. In the table, the scope of each data item within the set of entities is presented. For example, the plain form of the viewership vector is not visible to any of the entities of the protocol.

Table 5.2: The scope (columns) of the data items (rows).

Data items	Participants			
	TAM Service	TA	TV agents	
			All	L-selected
TV agents' viewership vector				
Plain form	✗	✗	✗	✗
Encrypted form	✗	✓$_{(1)}$	✗	✗
TAM result				
Plain form	✓$_{(2)}$	✗	✗	✗
Encrypted form	✓$_{(3)}$	✓$_{(4)}$	✗	✓$_{(5)}$

We are now ready to shown the main security features of the basic PrivTAM protocol. In general, the security features are inherited from the well known cryptographic primitives that are used in the protocol and we do not have to resort to formal methods and proofs [16].

- The TA cannot obtain information about the contents of the viewership vector (✓$_{(1)}$) and the encrypted TAM result (✓$_{(4)}$), since the ciphertexts are encrypted with the Paillier encryption.

  If for example, the TA stores a viewership vector (✓$_{(1)}$), its contents cannot be revealed unless all the participants in the threshold decryption (the L viewers and the TAM service) are malicious and collude towards this purpose.

- None of the viewers can manipulate the aggregate result, thanks to the viewership vector validation process (✓$_{(1)}$).

- None of L selected TV agents can obtain information about the content of the final TAM result (✓$_{(5)}$), since the collaboration of all of them and the TAM Service is required to reveal the TAM result in plain form.

- The TAM Service can obtain the content of the encrypted TAM result (✓$_{(3)}$) if and only all L agents send their partial decrypted results to it. The privacy of the viewers is preserved even if any number of the the L agents is malicious. However, if one or more of the L agents are malicious or simply fail, then the decryption of the aggregate result may not be completed. In this case, the computation of the TAM measurement has to be repeated with a fresh set of L agents for the threshold cryptosystem.

- At the end of the protocol, only the aggregate TAM result (✓$_{(2)}$) is revealed. Consequently, no viewer or strict subset of viewers can be associated with

any particular viewership vector, unless additional background information is given. That is, the protocol offers k-anonymity for $k = N$, against any external adversary, where N is the number of all the participants who take part in the measurement.

- The protocol is robust against malicious viewers. No viewer can manipulate the aggregate result or neither violate the privacy of other viewers. This holds even for large number of malicious users. If for example ℓ malicious viewers participate in a computation, the protocol assures $N - \ell$ anonymity for the legitimate viewers, even if all the malicious users collaborate.
- If necessary, one may extend the protocol with additional features. For example, if the PrivTAM computation should be protected from inaccurate result ($\checkmark_{(2)}$) reporting from the TAM Service, then multiple combiners could be introduced in Phase 3, in order to confirm the announced results.

5.5 Enhancement with Certified Demographic Data

In this section, we show how the described PrivTAM protocol could be enhanced to support privacy-preserving certified demographic data (Figure 5.4). This functionality is important for our protocol because we want the generated TAMs to provide certified aggregate demographic data. To achieve this goal, we utilize advanced cryptographic techniques, namely anonymous credentials [31], blind signatures [37, 30] and zero-knowledge protocols [69, 68].

Definition 5 (Anonymous credentials) *Anonymous credentials [31] allow users to prove possession of credentials without revealing any other information about themselves; when such a proof is carried out, it cannot be linked to previous uses of the same credential, or to any other identifying information about the user. Additionally, they give the users the ability to privately obtain credentials.*

Definition 6 (Blind Signature) *A blind signature [37, 30] is a form of digital signature in which the content of a message is disguised (blinded) before it is signed. The resulting blind signature can be publicly verified against the original, unblinded message in the manner of a regular digital signature.*

To support this new functionality in our proposed architecture (Figure 5.1) a new honest but curious authority needs to be added, the Certification Authority for Demographic Data (CA_d), which is responsible to validate and sign the viewers' demographic data. Additionally, the existing Trusted Authority (TA) needs to be

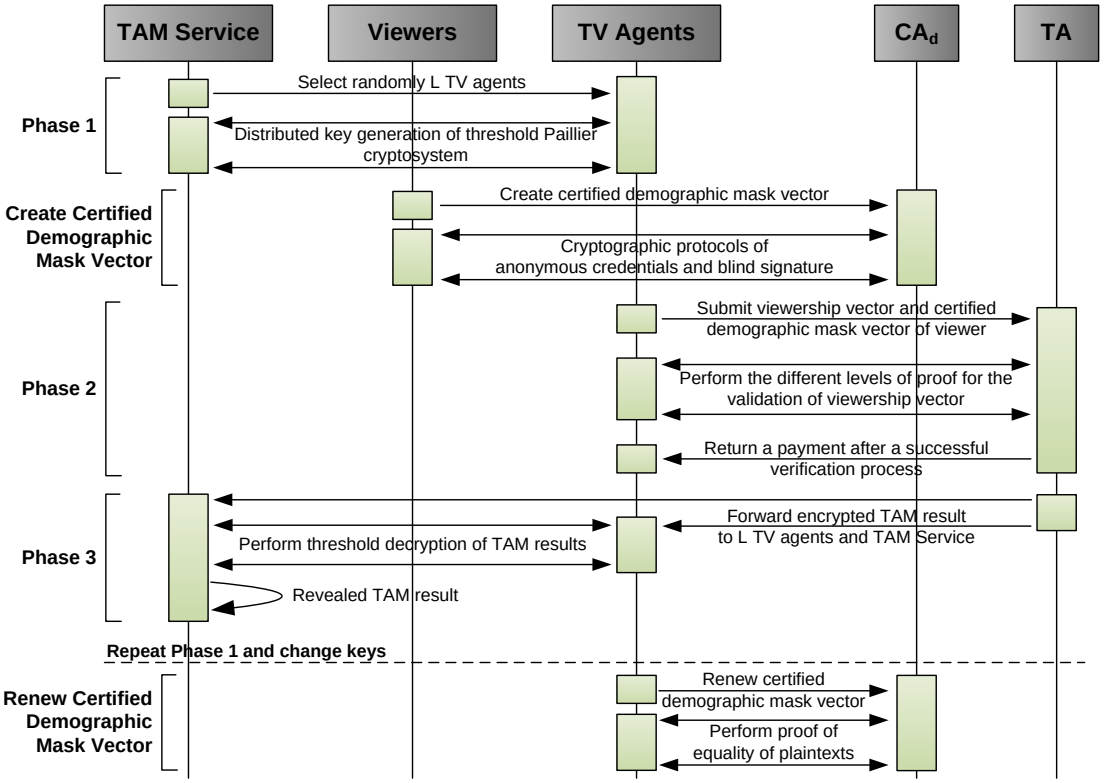

Figure 5.4: Interaction diagram of the PrivTAM system enhanced with certified demographic data.

enriched to support functionalities such as to validate the certified demographic data. These new features can be divided in three main parts:

1. *Creation of the certified demographic mask vector.* This first part consists from the protocol where the viewer proves (with anonymous credentials) to the CA_d the validity of their demographic data. If the proof is successful then a new protocol begins, wherein the viewer proves to the CA_d that the provided encrypted demographic vector corresponds to the viewer's proven demographic data. Finally, if this proof is successful, the CA_d signs (using a blind signature scheme [37, 30]) the encrypted demographic mask vector.

2. *Modifications in Phase 2.* This part consists of a modified protocol that is described in Phase 2 of Section 5.4.3. In this modified protocol, the TA validates the signature of the demographic mask vector and multiplies it with the viewership vector per channel. To support this functionality, the formatting of the viewership vector is altered compared to the previous protocol in Phase 2.

3. *Certified demographic mask vector renewal.* Finally, the last part consists of the protocol where the viewer proves to the CA_d that a new encrypted demographic mask vector (using the changed public key for the computation) encrypts the same demographic data with the previous signed one.

The full descriptions of these three parts are given in the following subsections.

5.5.1 Creation of the Certified Demographic Mask Vector

The creation of a new certified demographic mask vector is a process that requires two main steps:

- The first step is for the viewer to prove to the CA_d, using anonymous credentials, the validity of their demographic data, i.e. that she is a female at an age between 25 and 40 years old.
- The second step is to prove to the CA_d that the encrypted demographic mask vector, which the viewer has prepared with the current Paillier public key, encrypts the same demographic data. If the proof is successful, the CA_d will sign the demographic vector of the viewer. This has to be achieved without revealing the final encrypted demographic vector that will be provided by the viewer to the TA.

In more details, in the first step the viewer should prove their attributes to the CA_d and the CA_d should have the capability to identify a demographic mask vector has been signed before for this viewer. This is required because every viewer should have only one valid demographic mask vector at each moment in time. The proven viewer attributes could be, the gender (male or female), the age group (≤ 24, $25-40$, $41-55$ or > 55) and educational level of the viewer. If this process is successful then the second step begins.

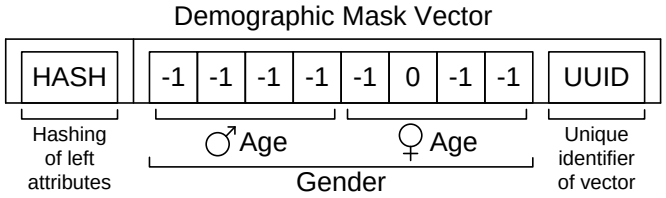

Figure 5.5: Example of a demographic mask vector.

In the second step, the demographic mask vector is generated by using the viewer's demographic data. This demographic mask vector, as is shown in Figure 5.5, consists for three parts. The first part of left, the UUID, is a random

ID (a nonce, i.e. a number used once) and is representative of the demographic vector. For this reason, we use a universally unique identifier (UUID) where the probability of collisions is practically negligible. The central part is the encrypted representation of demographic data where the number "0" corresponds to the viewer's demographic data and the number "-1" corresponds to the opposite. In Section 5.5.2, it is shown why this form (0 and -1) is required. For example, in our case the demographic vector in Figure 5.5 shows that the gender is female and the age is between 25 and 40 years old. The last part on the left, the HASH, is a hashing of the encrypted representation of demographic data and the UUID. The cryptographic hash function that used for this purpose is the SHA-2 (SHA-256 or SHA-512). The protocol that is running between the viewer and the CA_d is:

1. The viewer generates D demographic mask vectors. Each of them has different UUID and ciphertexts that encrypt the same data. Consequently, the HASH value is also different for every vector.

2. For every demographic vector, the viewer blinds the HASH h by using blind RSA signature scheme. We suppose the CA_d has RSA key pair where the public key of the modulus n_r and the public exponent e and the private key consists of the modulus n_r and the private exponent d. The blinded HASH h' is calculated by using a blinding factor $r^e \bmod n_r$ as:

$$h' = hr^e \bmod n_r \tag{5.2}$$

   where r is a random value, such that r is relatively prime to n_r (i.e. $gcd(r, n_r) = 1$). After that, the viewer only sends to CA_d the D blinded hashes h'.

3. The CA_d randomly selects $D-1$ blinded hashes h' and asks from the viewer to send all the required values with the purpose to regenerate the selected blinded hashes.

4. The viewer sends to the CA_d the values: the UUIDs, the random numbers of the encrypted demographic data and the random values r of the blinding factors.

5. If the CA_d successfully regenerates the $D-1$ blinded hashes h' then the blinded signature s' of unrevealed blinded hash is calculated as:

$$s' = (h')^d \bmod n_r \tag{5.3}$$

   and sent to the viewer.

6. Finally, the viewer removes the blinding factor to reveal s as:

$$s = s'r^{-1} \bmod n_r \tag{5.4}$$

where the signature s is now valid for the hash h and as a result for the demographic mask vector.

5.5.2 Modified Phase 2 with Certified Demographic Data

To support the added functionality for the certified demographic mask vectors of viewers, modifications need to be made to Phase 2. The changes that are required do not notably alter the main aspects of the computational task of the existing process. The form of the viewership vector is altered, such that all the values in the vector that do not represent the viewer's demographic data are 1, whereas the value in the vector representing the viewer's demographic data is 1 for the channel the viewer is watching and 0 for the rest of the channels. This new form is shown in Figure 5.6 as viewership vector. The TA performs the next two simple steps:

- First, the TA generates and checks if the hashing of the UUID and the ciphertexts of the demographic data is equal with the hash of the demographic mask vector.
- Second, the TA checks if the hash h is equal to $h = s^e \bmod n_r$. This shows that the signature is valid.

The UUID is used by the TA to identify the demographic vector and it is simultaneously used for the authentication of the TV agent to verify that the viewer has not taken part again in the current TAM. After the successful validation of the demographic mask vector, the viewer should prove to the TA that:

1. Every ciphertext in the viewership vector lies in the set $S = \{0, 1\}$.
2. Next, the TA multiplies every channel of the viewership vector with the certified demographic mask vector, as is shown in Figure 5.6, and the product is the validated viewership vector. So, every ciphertext in the validated viewership vector should also lie in the set $S = \{0, 1\}$.
3. The multiplication of the ciphertexts in every channel of the validated viewership vector should lie in the set $S = \{0, 1\}$.
4. Finally, the multiplication of all ciphertexts in the validated viewership vector should equal to "1". This means that the viewership data are not null.

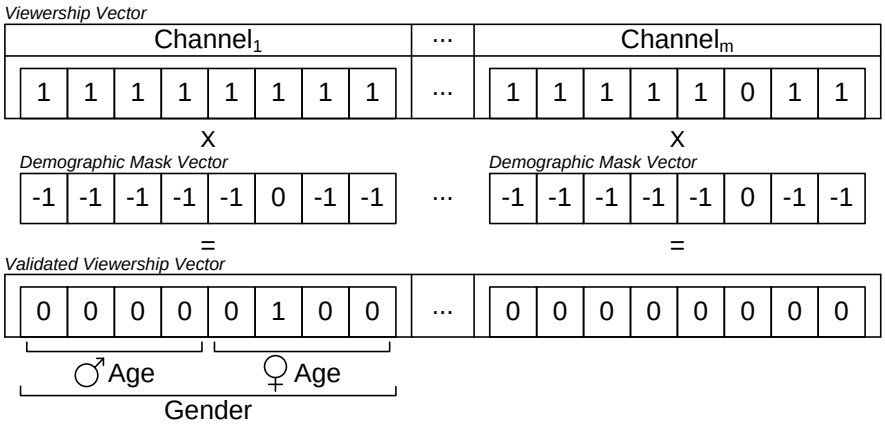

Figure 5.6: Example of the modified viewership vector.

5.5.3 Certified Demographic Mask Vector Renewal

In this section, we suggest a method to renew the certified demographic mask vector that is easier and more efficient than the creation phase described in Section 5.5.1 This process requires that there is an existing certified demographic mask vector. The renew process is necessary to be repeated whenever the Paillier keys are changed for the measurement process of TAMs. The steps that the viewers should follow to renew their certified demographic mask vectors are:

1. The viewer creates a new demographic mask vector that consists from the UUID of the old mask vector, the demographic data encrypted with the new Paillier public key and a new hashing of the previous demographic data. This new demographic vector and the previous certified one are sent to CA_d.

2. A zero-knowledge proof of equality of plaintexts [18] begins between the viewer and the CA_d, that will prove to the CA_d that the two encrypted demographic data are equal without revealing anything about the plaintext of the demographic data. This proof is described in detail in Section 5.5.4.

3. If this verification process is successful, the CA_d will check that the UUID of the new demographic mask vector is the same as the old one and that the hashing is correct, and then it will sign the hash h with the RSA private key d as follows: $s = h^d \bmod n_r$.

This proposed demographic mask vector renewal process does not need to use blind signatures (like the creation process) because we do not have linkability between the two certification processes.

5.5.4 Security Discussion for the enhanced PrivTAM protocol

We will show that the enhanced protocol fulfills the requirement for certified demographic data of the viewers in addition to the features of the basic PrivTAM protocol. The critical data of the computation are again the viewership records and the demographic data of the viewers, but there is an additional entity and the protocol makes use of extra cryptographic tools with respect to the basic protocol. In particular, the entities of the enhanced protocol are the viewers, the TA, the TAM service, and the certification authority for the demographic data (CA_d). Moreover, the enhanced protocol presupposes the use of an anonymous credentials infrastructure. The same security model assumptions hold for the TA (HBC), the TAM service (HBC) and the viewers (can be malicious). The new entity CA_d is also assumed to be HBC.

All the facts of the basic protocol hold for the enhanced protocol too. In addition, in the enhanced protocol a malicious user cannot misreport its demographic data. Anonymoys credentials, blind signatures and zero knowledge proofs are used to assure that only the right (i.e. corresponding to the right demographic data class) cells of the viewership vector can contain non-zero values. Finally, the privacy of the viewers is preserved also in the renewal process, because the CA_d cannot see the unique UUID of the demographic mask vectors. Thus, the CA_d entity cannot threaten the privacy of the viewers, even if it would conspire with the TA. The additional cryptographic tools are anonymous credentials [31] and blind signatures [37, 30], with security proofs given in [15] and [83], respectively. Finally, the description of the zero-knowledge proof for the equality of two plaintexts is given below, whereas its security analysis can be found in [18].

Zero-Knowledge Proof of equality of plaintexts [18]. Let $n_{p_1}, \ldots, n_{p_w}$ be w ℓ-bit RSA modulus. Given w encryptions $c_j = {g_j}^m {r_j}^{n_{p_j}} \mod n_{p_j}^2$ under the assumption that the decryptions of the c_j lie in an interval $[0, 2^\lambda[$, the prover P convinces the verifier V that the c_j's encrypt the same message m.

1. P picks at random $\rho \in [0, 2^\ell[$ and $s_j \in \mathbb{Z}^*_{n_{p_j}}$ for each $j \in \{1, \ldots, w\}$. Then he computes $u_j = {g_j}^\rho {s_j}^{n_{p_j}} \mod n_{p_j}^2$ and commits the u_j.

2. V chooses a random challenge e in $[0, A[$ and sends it to P.

3. P computes $z = \rho + me$ and $v_j = s_j {r_j}^e \mod n_{p_j}$ and sends z and the v_j for each $j \in \{1, \ldots, w\}$.

4. V checks that $z \in [0, 2^\ell[$ and that ${g_j}^z {v_j}^{n_{p_j}} = u_j {c_j}^e \mod n_{p_j}^2$ for each $j \in \{1, \ldots, w\}$.

According to Theorem 3 of [18], it holds that T iterations of the above protocol is a statistical zero-knowledge proof of membership (against an honest verifier) that data $\{c_1, \ldots, c_w\}$ encrypt the same λ-bit message, for any non-zero parameters A, T and λ such that $1/A^T$ and $2^{\lambda-\ell}A$ are negligible.

Table 5.3: The scope (columns) of the extra data items (rows) in Section 5.5.

Data items	Participants				
	TAM Service	CA_d	TA	TV agents	
				All	L-selected
Viewer demographic data					
Anonymous demographic data	✗	✓$_{(6)}$	✗	✗	✗
Demographic mask vector					
– Creation process	✗	✗	✗	✗	✗
– Modified Phase 2	✗	✗	✓$_{(7)}$	✗	✗
– Renewal process	✗	✓$_{(8)}$	✗	✗	✗

The main extra security features of the enhanced protocol with respect to the basic protocol are (Table 5.3) are:

- The CA_d cannot link the credentials (✓$_{(6)}$) of a user with a particular demographic mask vector (✓$_{(7)}$ or (✓$_{(8)}$), since the CA_d gets access and signs only the blinded hash h' of the demographic mask vector.
- The digital signature s of the demographic mask vector (✓$_{(7)}$) is not linkable with its blinded version s'. Thus, the demographic mask vector cannot be linked to the demographic data even if the TA would collaborate with the CA_d.
- The modified version of Phase 2 (Section 5.5.2) ensures that the validated viewership vector contains certified demographic data and the final aggregate TAM result is also certified with respect to the demographic data.
- The usage of the certified demographic mask vectors (✓$_{(7)}$) in Phase 2 does not leak any information about the participating users.
- After the renewal process of an already certified demographic mask vector (✓$_{(7)}$), the CA_d is unable to link the new demographic mask vector with the demographic data of the viewer. Although, the current and the new demographic vectors are linked (✓$_{(7)}$ and ✓$_{(8)}$), no information is leaked about the demographic data of the viewers. This fact preserves the unlinkability between the demographic data (✓$_{(6)}$) and the encrypted demographic mask vector.

5.6 Experimental Evaluation

To evaluate our approach, we initially developed a PrivTAM prototype that implements the main privacy-preserving protocol for calculating a TAM. Furthermore, we performed a series of experiments to evaluate the computational requirements of a TA server. Additionally, basic implementations for the components and functionality for the enhancements proposed in Section 5.5 have been developed, namely for blind signatures, zero knowledge proofs, anonymous credentials and micropayments and preliminary results indicate the feasibility of the approach.

5.6.1 The PrivTAM prototype

The PrivTAM prototype can be separated into two main parts, the first being the application on the Smart TVs and the second the application on the TA. The Smart TV application is implemented using the Google TV platform[1] and the Java for Android SDK. Using other Smart TV platforms, it also possible since the only requirement is for the platform to support operations with big integers. The application on the TA is also implemented in Java. Both applications use the cryptographic primitives of the Paillier Threshold Encryption Toolbox [137]. In this library, a centralised mechanism (with a trusted dealer) for threshold key generation [43] is implemented, instead of a distributed Paillier key generation [108]. In our view, this is enough for this prototype implementation. The TV agents use production-ready cryptographic libraries and employ 1024 bits RSA X.509 certificates.

We performed an experiment of the PrivTAM calculation, where 6 TV agents, the TA and the TAM Service participated and four channels exist. Each agent generated random values for the submitted viewership vector, as well as for the gender and the age of the viewer. Initially, the TAM Service randomly chooses two of the participating TV agents ($L = 2$), $TVAgent_2$ and $TVAgent_5$, for the first phase of the protocol. Therefore, the final encrypted measurement will be decrypted by $TVAgent_2$, $TVAgent_5$ and the TAM Service (n, $t = 3$ parties).

Next, each TV agent encrypts the viewership vector and transmits it to the TA for validity check. This process in our experiments takes less than 8 seconds. Once the viewership vector is checked, the TA multiplies it with the current encrypted TAM result. In Table 5.4 the values used to create the viewership vector of each agent are shown, along with the resulting current encrypted measurement after the submitted viewership vector is calculated by the TA.

At the end of the computation, the TA sends the encrypted results to $TVAgent_2$,

[1] http://www.google.com/tv/

Table 5.4: Example of a PrivTAM.

TV Agents Values				**Current Encrypted TAM**			
Agent	Channel	Gender	Age	$Channel_1$	$Channel_2$	$Channel_3$	$Channel_4$
$TV\,Agent_1$	$Channel_3$	Male	23	0000 0000	0000 0000	1000 0000	0000 0000
$TV\,Agent_6$	$Channel_1$	Female	45	0000 0010	0000 0000	1000 0000	0000 0000
$TV\,Agent_2$	$Channel_1$	Male	32	0100 0010	0000 0000	1000 0000	0000 0000
$TV\,Agent_4$	$Channel_4$	Female	29	0100 0010	0000 0000	1000 0000	0000 0100
$TV\,Agent_3$	$Channel_3$	Female	53	0100 0010	0000 0000	1000 0010	0000 0100
$TV\,Agent_5$	$Channel_3$	Female	22	0100 0010	0000 0000	1000 1010	0000 0100

$TV\,Agent_5$ and the TAM Service. The TAM Service collects the partial decryption results from $TV\,Agent_2$ and $TV\,Agent_5$, and combines the partial decryption results. The decrypted TAM result is shown in the last row of Table 5.4, where $Channel_3$ has the highest audience (50%) and the 66.6% of the viewers were women. A snapshot of the application during the execution of the experiment is shown in Figure 5.7.

Figure 5.7: A snapshot of the $TV\,Agent_5$.

5.6.2 Computational performance evaluation

To evaluate the computational requirements for a TA Server of the PrivTAM system, we performed a series of experiments where we calculated the overall execution time on the server for different number of clients and the execution time per client for different number of threads on the server. In this series of experiments, we used as parameters 4 TV channels with 4 age groups and 2 gender categories. For our experiments, we used a simple workstation with a CPU Intel Core i5 750 (4 cores) at 2.67 GHz and 4 GB RAM as the TA Server. The clients were simultaneously executed in a HP Proliant DL370 G6 server with two CPU Intel Xeon X5650 (6 cores) at 2.67 GHz and 24 GB RAM. The server and the clients were implemented in Java and were running in 64-bit and 32-bit Java VMs. The encryption of the viewership vectors and the creation of proofs are offline processes and they are not included in the execution times.

In Figure 5.8, the overall execution times of the verification process at the TA for different number of viewers are shown. From this graph, we can see that the execution time is linear relative to the number of viewers that take part simultaneously in the computation. Additionally, as shown in the graph, the 64-bit Java was faster than the 32-bit Java and the execution time was almost four times lower.

Furthermore, the execution times per client for a different number of threads on the server are shown in Figure 5.9. As shown in the graph when the number of threads on the server is greater than the number of cores, the execution time is almost equal to 728 msec and 2970 msec for 64-bit and 32-bit Java, respectively. This happens because the server takes advantage of 100% of CPU usage. Thus, if we want to calculate a TAM (with the same parameters as in this experiment) within 10 minutes and if we have 20,000 participants with 64-bit Java, then 24 workstations are required to achieve this goal.

The proposed architecture of the PrivTAM system does not require any dedicated hardware to be installed and managed at the viewers' side. Therefore, the cost of maintaining such a system is expected to be lower than the current practice. Both solutions include a server that conducts the measurement computation, but the cost of supplying and maintaining a dedicated metering device for each of the participants in the case of the People Meter increases significantly the cost of the system operation. In the case of PrivTAM deployment is easier, as users only need to activate the PrivTAM application on their SmartTV, no dedicated equipment needs to be installed at the participants' side for the measurement to take place. Additionally, the PrivTAM solution is scalable, supporting much larger numbers of participants, only adding cost to the processing power of the server, while the metering device solution increases the operation costs for each additional participant.

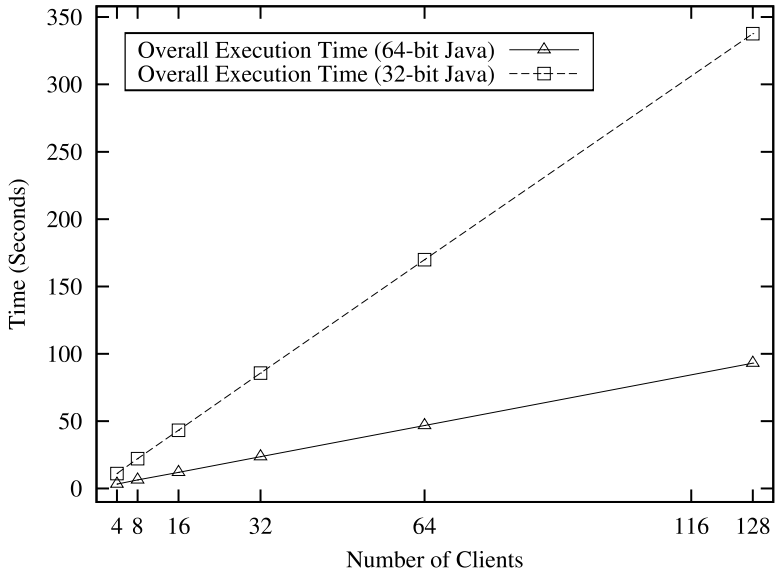

Figure 5.8: Overall execution times on server for different number of clients.

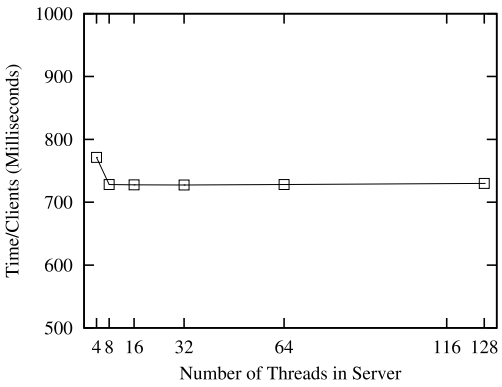

(a) Execution Time per Client for 64-bit Java.

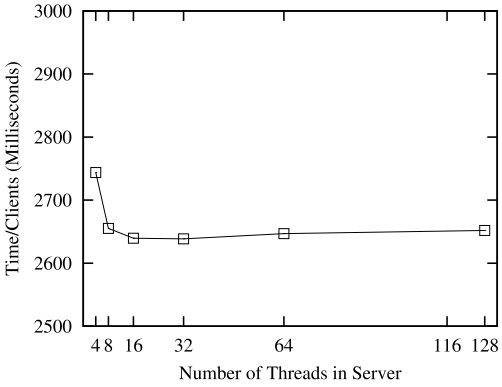

(b) Execution Time per Client for 32-bit Java.

Figure 5.9: Execution times per client for different number of threads on server.

Furthermore, we investigated which step of the verification process is more computationally expensive for the server and the clients. For this purpose we used Java profiling techniques and particularly the tool VisualVM. From this analysis, we found that for server more than 90% of the execution time is used by the 4^{th} step (proof in Section 5.4.4) of the vector validation process. The same happens at the viewers' side for the offline process where the encryption of viewership vector and the creation of proofs (1^{st} step of the proof in Section 5.4.4) take place.

5.7 Discussion

The introduction of Internet connectivity and computation capabilities to contemporary television systems, opens the possibility of conducting TAM measurements using larger samples of viewers. In this work we propose a new approach for computing privacy-preserving TAMs and test the applicability of the proposed solution.

Enhancements to the initial measurement protocol offer accountability for the participants' demographic data, leading to more reliable TAM ratings. The accuracy and trustworthiness of the produced results act as strong incentives for TAM Services to adopt the PrivTAM system.

From the viewers' perspective, PrivTAM offers the privacy assurance necessary for them to participate in a TAM system, while fair compensation can be offered for their participation, returning some of the economic benefits of TAMs back to the viewer.

Furthermore, PrivTAM can support on demand alternative kinds of TAM measurements, providing interesting information about audiences to the TV industry. These results are achieved without using any specialised equipment and can take into account data from multiple broadcast sources. Overall, we believe this is an interesting approach, which, using contemporary cryptographic tools, achieves reliable TAM ratings, while protecting the viewers' privacy.

A future direction for the PrivTAM approach would be to investigate the possibility to use a decentralised architecture, such as a peer-to-peer topology, where the TV agents would self-organise, compute amongst themselves and validate the correctness of the TAM result. In this case, a TA would not be required. Finally, an interesting extension of the PrivTAM would be to support polls where viewers could express their opinions about current social or political issues.

Chapter 6

A Personal Data Management Portfolio for Privacy-preserving Consumer Data Utilization

6.1 Introduction

During electronic and physical transactions today, users are required to provide their personal information, without being offered the option to carry on with their purchase anonymously. In most e-purchases customers are requested to register, creating a profile with the company, upon which they have little or no control. Even with physical transactions, it is often required that customers register their information with the shop in order to complete a purchase. Some shops offer membership cards that entitle customers to member privileges and offers. These shopping profiles provide valuable information to interested parties, but can be a significant source of private information leaks.

Processing of consumer data is a popular method of extracting valuable information to serve commercial interests. Patterns of consumer behavior and habits can provide valuable insights to companies on how to promote their products and handle their customer relationships. Loyalty cards are a common way for companies to acquire customer transaction data in exchange for personalised prices and offers. These profiles often contain customer Personal Identifiable Information (PII), such as their name, address, zip code and phone number and can be used to extract sensitive information about a person's attributes, habits and beliefs.

In this work we propose that individuals store and manage their transaction data in a personal portfolio, which remains under their control. A searchable encrypted version of the transaction data, contained in the Portfolio, is outsourced to an external cloud storage server, providing controlled access to accountable

Portfolio data for interested data consumers in a privacy-preserving way. The proposed solution offers a new paradigm regarding transaction data management systems compared to the existing practice, offering enhanced utilization capabilities and accountability assurances for the transaction data, while the data owners retain control of their information

Giving control of the data back to the user raises accountability issues on the accuracy of the data the user reveals to the data consumer. Users should not be able include false information in their profiles to deceive data consumers. By combining contemporary advanced cryptographic techniques to create appropriate transaction protocols, the proposed system provides information accountability while preserving the owner's privacy. The Portfolio functionality is supported by an agent-based system where software agents representing users exchange Portfolio information according to the users' preferences.

Apart from keeping consumer profiles, the Portfolio approach can be also be used in other applications, such as tax records, health records, biographic or even opinion records for statistics and polls. Although each application entails different challenges and requirements, the general requirements remain unchanged:

- Protecting the individual's private information.
- Providing an efficient and reliable method of accessing the Portfolio data for legitimate users.
- Offering data accountability, protecting data processing entities from ill behaving and malicious users.

In this work we utilize contemporary advanced cryptographic techniques to propose a personal data Portfolio that is used for the benefit of the individual. The proposed idea can be implemented using existing hardware and technologies, leading to applications that can enhance user privacy protection in everyday practice. This work, which is a follow up work of [132], significantly extends and improves the construction of the Portfolio and updates the incorporated cryptographic building blocks with more recent developments from the field of searchable encryption. The resulting new architecture is more concise and the corresponding protocols that achieve the desired functionality and attributes are more straightforward.

6.2 Portfolio Concepts and Architecture

The main idea of the Portfolio approach is the following: *A user's Portfolio contains the user's transaction history. The Portfolio is stored at the owner's*

side and a searchable encrypted version of the Portfolio transaction dataset is outsourced to an external cloud storage server, to support access to its contents under well-defined rules.

The focus of this work is primarily on privacy-preserving transaction data management, however some identity management issues also arise. Although individuals use anonymous credentials to perform their transactions, it is possible to support the storage of some of the individual's PII in the Portfolio as well. Depending on the application, this information can be used to share some demographic information along with the transaction data (e.g. age group, gender).

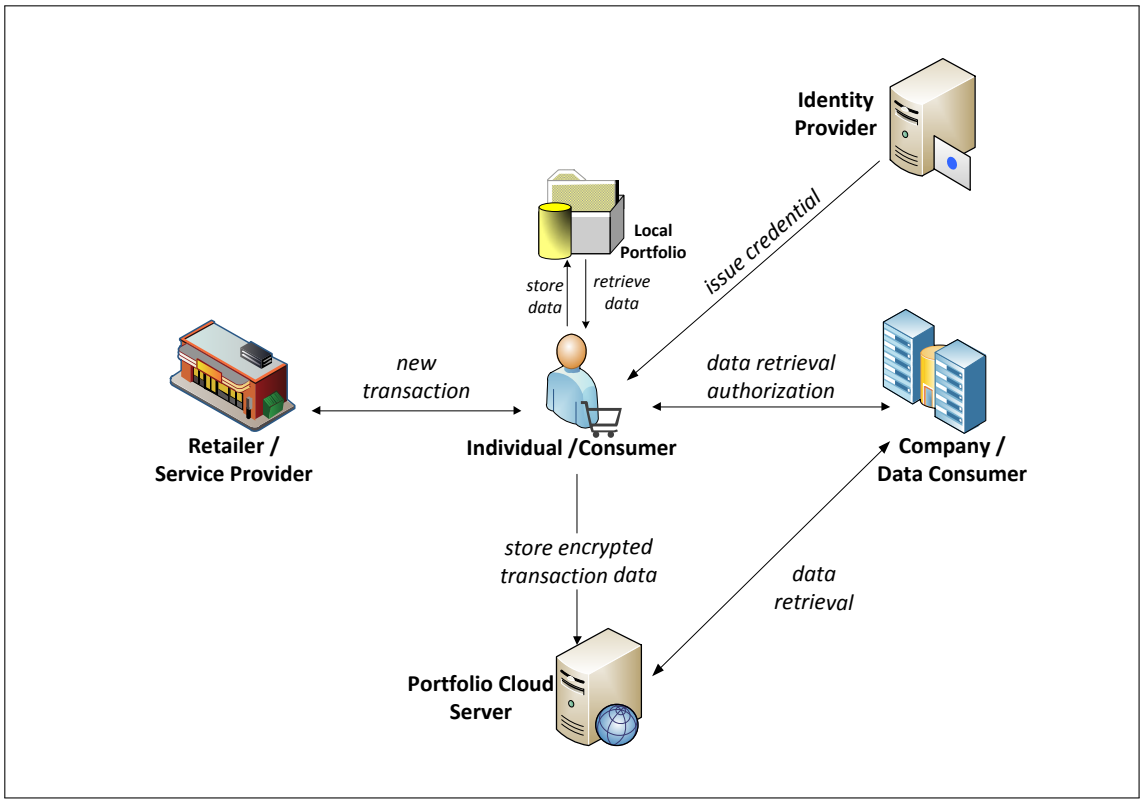

Figure 6.1: The Portfolio architecture

The entities that interact in the Portfolio system are the following:

- The individual, or user, who owns a Portfolio and acts as a customer during transactions.
- Retailers and service providers that the individuals perform transactions with, e.g. super markets and e-shops.

- A cloud storage service that supports online storage of the encrypted Portfolio transaction dataset.

- Companies that act as data consumers and cloud service clients, interested in acquiring data from a users' Portfolios.

- An identity provider that issues credentials to the Portfolio users.

The identity provider is considered a trusted third party. The cloud storage server is a semi-trusted server which acts as honest-but-curious, and additionally does not collude with clients. Software agents, i.e. software that acts on behalf of the user, manage the Portfolios. This means that they act as always-on local services which represent the users and allow for the communications and the data exchanges between the participants.

Naturally, a retailer can also be a data consumer, interested in acquiring access to Portfolio data either to provide personalised services and offers to customers or to analyze consumers' behavior to determine correlations between products and adapt its marketing strategies accordingly. Nevertheless, in the general Portfolio architecture, the retailer and the data consumer are represented as different entities. There can be retailers that are not interested in consumer data processing, however consumers would still import their transactions with them into their Portfolio. Additionally, there can be data processing companies that do not perform any transactions directly with the individuals and are only interested in transaction data mining.

The critical resource within the Portfolio system is the transaction (T_i) between an individual and a retailer, which is stored into the individual's Portfolio, where i is the transaction identifier in the Portfolio database. The transaction (T_i) also constitutes the system's information unit. An individual's Portfolio contains and manages a set of transactions and a transaction is the smallest unit of of information exchanged between the Portfolio and an authorised data consumer.

The Portfolio system architecture is presented in Fig. 7.1, where the Portfolio entities and their interactions are illustrated. The two main operations of the Portfolio are data insertion during new transactions and data retrieval by data consumers. Transaction data is created during new transactions between a retailer and a customer, and is stored into the customer's Portfolio. Transaction data is retrieved by a data consumer agent from an individual's Portfolio cloud storage service. These operations and their protocols are described and analyzed in Section 6.4.3.

6.2.1 Portfolio Use Case

In this section, we describe a Portfolio use case to better illustrate its possible uses and explain its intended functionality. In general, many similar scenarios can be supported, concerning different domains and applications.

Suppose an individual is interested in purchasing some items from an online retailer, for example groceries. Depending on the prospective customer's consumer profile, the retailer will make an offer, in order to attract the customer's business. To do so, the retailer will want to access the customer's purchase history, to establish the amount of product the customer buys per month or year and the items the customer usually buys along with groceries.

In today's practice, this kind of information about the customer is gathered by the retailer's profiling process. Customers register with the retailer and profiles are maintained and stored at the retailer's side, containing the customers' purchase history, along with their personal identifiable information (such as phone number, address, etc).

This practice has some important drawbacks, both regarding the consumer's privacy, and the retailer's ability to make accurate personalised offers. First, the profile's protection becomes the retailer's responsibility, which is a financial burden. Additionally, even if the retailer means well, data leaks can occur and if revealed, the retailer's reputation is severely affected. From the consumer's side, the profile data are offered to the retailer virtually condition-free and the ways they are used are unknown to their owner, unless some major leak is revealed and the consequences become apparent (identity theft, credit card theft, etc.). Apart from these drawbacks, today's customer profiles have another important limitation. They only contain information on the customer's purchase habits with the particular retailer. The consumer may have a broader purchase profile that could be utilized for the customer's (and the retailer's) benefit.

The Portfolio approach allows the individual to maintain and control their consumer profile, preserving their right to control access to it and know about the uses of its data. Additionally, the Portfolio profile contains a more inclusive purchase history, allowing the user to better utilize their consumer profile and the retailers to make more accurate personalised offers to their customers. The accountability of the Portfolio data is ensured, so that customers cannot fabricate fake transactions or tamper with the existing transaction contents.

6.2.2 Requirements

Before describing the protocols for the portfolio operations, we identify their functional, privacy and accountability requirements. These requirements are addressed later on by utilizing and combining appropriate cryptographic techniques

(analyzed in Section 7.4.1).

The identified *privacy and accountability requirements* for the portfolio operations are the following:

- Transaction information can be disclosed to data consumers in a privacy-preserving and controlled way.
- The Portfolio data integrity and accountability needs to be ensured. This requirement can be subdivided into three aspects/sub-requirements:
  - The Portfolio owner cannot tamper with the contents of a transaction.
  - The Portfolio can contain only transactions performed by its owner, transactions from other consumers cannot be added to the Portfolio.
  - Fake or falsified transactions cannot be added to the Portfolio.

  We do not consider the mandatory inclusion of all new transactions into the Portfolio a requirement. We consider it an individual's right to omit transactions from their Portfolio, or even maintain multiple different Portfolios, for different categories of transactions. Individuals have incentive to include transactions in their Portfolio(s) to enrich their consumer profiles, nevertheless they should be able to omit or group them according to their preference, as long as the transactions included in their Portfolio(s) are valid, tamper-proof and have been performed by the Portfolio owner.

  However, it is straightforward to support the mandatory inclusion of all transactions in the Portfolio for applications that require it (e.g. tax records), under the assumption that the consumer does not withhold the fact that they own a Portfolio. In such cases this assumption could be externally enforced, for example by the government in the case of tax applications, making the use of the Portfolio obligatory in all transactions.
- The cloud storage service cannot determine the identity of the Portfolio owner, the contents of the Portfolio or the contents of the search queries.
- Retailers do not keep any other information about their customers, apart from anonymous transaction information.

The *functional requirements* of the portfolio operations are the following:

- The Portfolio needs to support privacy-preserving search operations on the transaction data.

- Efficient additions of new transaction information to the Portfolio should be supported.
- The results returned regarding a company's query, should contain all the relevant transactions included in the up-to-date version of the Portfolio.
- Data exchange operations need to be efficient, to ensure the applicability of the solution.

6.3 Portfolio Building Blocks

To address the requirements identified in Section 6.2.2 we design the Portfolio functionality, containing appropriate privacy-preserving protocols, based on cryptographic primitives. These primitives are used as building blocks, whose attributes, appropriately combined, fulfill the Portfolio requirements. In the following paragraphs, high-level descriptions of the cryptographic building blocks used to construct the Portfolio operations and protocols are provided, focusing on their functionality and their attributes.

6.3.1 Private Credentials

Private credentials allow users to acquire credentials and demonstrate them without revealing their identity. Using private credentials, individuals can use different unlinkable pseudonyms with each service provider they interact with, based on the same credential issued for their Portfolio by the identity provider. The private credential system can also provide certified attributes by the identity provider, for the individual to selectively reveal attributes (e.g. their age range, based on their date of birth). Anonymity revocation is also supported by the private credential system, allowing identities of ill-behaving users to be revealed under well-specified conditions. A private credential system, apart from users who receive credentials and the certification authority that verifies user public-private key pairs, comprises of organisations which grant user credentials, verifiers who verify them, and an anonymity revocation manager that allows revocable anonymity.

We use private credentials as a "black box" component in the portfolio architecture, therefore we will not go into the details of its functionality. A more detailed description can be found in (Section 2.3.1).

The entities participating in an anonymous credential system are individuals, companies and trusted third parties (e.g. government services), which can assume the roles of issuers, recipients, provers and verifiers. A credential is created by an issuer for a recipient by executing the issuing protocol. The recipient (i.e. the

credential owner) can then create a credential proof, to be used by a verifier to verify the validity of the credential (proving protocol).

To be able to issue anonymous credentials, an issuer needs to generate a public key pair and create specifications of the structures of the credentials issued. These specifications and the issuer's public key are then published to be used during the proof protocol.

In order to acquire an anonymous credential, a user chooses a master secret key, according to the agreed upon system parameters (bit length, groups to be used). This secret key enables the creation of multiple unlinkable pseudonyms by the user, to be used with different service providers. The issued credential consists of the issuer's public key, the credential structure (necessary for the verification of its validity) and the attribute values.

During the proving protocol, the prover (i.e. the credential owner) creates a proof on behalf of the verifier that proves ownership of a certain credential. The verifier checks the validity of the given proof. Credential attribute values contained in the proof may or may not be revealed to the prover, according to the settings of the proof creation process.

6.3.2 Searchable Encryption

Searchable encryption solutions provide the ability to store information in an encrypted form, yet allow search operations to be performed on the dataset, in order for authorised users to be able to locate the information they are interested in, without needing to decrypt the whole dataset. More details on searchable encryption are provided in Section 2.3.2.

Within the Portfolio context, a searchable encryption scheme is needed that fulfills the following requirements:

- Updatability of records. New transaction records need to be efficiently added to the encrypted dataset.
- Multi-party functionality. The Portfolio is maintained by the data owner, yet multiple (authorised) users can perform searches on the dataset.
- Query expressiveness. It is desirable (yet not mandatory) that more complex than single-word queries are supported.

Searchable Symmetric Encryption (SSE) is a practical method for searchable encryption that provides a reasonable trade-off between efficiency, functionality and security. The SSE functionality is described in more detail in Section 2.3.2.

For the purpose of our work we consider the OXT symmetric searchable encryption scheme [35, 80, 34] more appropriate and closer to our requirements.

However, as research on the field of searchable encryption moves forward, the searchable encryption scheme used within the Portfolio can be updated accordingly. In the following paragraphs we provide a high-level description of the OXT scheme characteristics and attributes, which are useful in the Portfolio context.

The dynamic multi-client OXT scheme. The selected symmetric searchable encryption scheme consists of the basic OXT scheme [35], appropriately enhanced to support multi-client search functionality [80], and dynamic updateable encrypted datasets [34]. Additionally, the OXT scheme supports conjunctive search and general boolean queries and remains efficient for very large databases. More information on the dynamic multi-client OXT scheme can be found in Section 2.3.2.

In brief, in the basic OXT protocol the server holds encrypted pointers to documents in a dictionary D and the search client (and data owner) holds a list of keywords. The protocol output is the set of encrypted pointers to the documents containing the client's keywords. To retrieve the documents, the client decrypts the pointers and obtains the matching (encrypted) documents. The server does not perform the decryption and cannot learn the keywords in the client's query. The OXT scheme computational complexity is independent of the number of documents in the dataset and scales with the number of documents matching the least frequent keyword in the search query.

To support multi-client search functionality, the basic OXT scheme is enhanced so that the data owner outsources an encrypted dataset to an external server and allows other parties to perform queries on the encrypted data by providing them with search tokens for specific queries. In order to achieve that, the data owner provides the search-client with a set of trapdoors (determined by the query and independent of the searched data), that can be transformed into search tokens for the OXT scheme. To ensure that the searcher is authorised by the data owner, the trapdoors provided by the data owner are homomorphically signed and so are the transformed search tokens submitted to the server. These tokens and signatures are unforgeable even by fully malicious clients [80]. Homomorphic signatures enable a signer to create a signature on a document and allow other parties to make predefined alterations to it and obtain a new signature on the altered document without interaction with the signer [82, 13].

To support dynamic additions to the encrypted dataset after it has been uploaded, the dataset owner needs to be able to compute the labels for the new data, to be added to the dictionary by the server. The scheme remains the same as the static OXT scheme, with the addition of an Update protocol and an auxiliary encrypted database EDB^+ and dictionary D^+ are used, which are initially empty and change as updates happen. Searches in the dynamic databases are performed by the server by first searching the dictionary D for keyword w, as in the static

case, then re-computing all the labels corresponding to w in D^+. For this operation the data owner needs to provide the value of the keyword-specific counter, therefore a dictionary D_{count} is maintained, associating each added keyword with its counter value.

6.4 Portfolio Functionality and Privacy-preserving Protocols

In this section we present the proposed solution for privacy-preserving and accountable transaction data management. We first provide an overview of the Portfolio functionality, we proceed with some setup functions and assumptions, and then we present the main protocols of the Portfolio functionality.

6.4.1 Functionality Overview

Individuals use anonymous credentials for their interactions with the retailers and the Portfolio cloud server. After acquiring a credential from the Identity Provider, they are able to issue multiple unlinkable pseudonyms to be used with different retailers. Additionally, the Portfolio owner can prove that two pseudonyms belong to the same credential. This feature can be utilized to enhance the accountability of the data insertion protocol (Section 6.4.3) and ensure that the submitted transactions belong to the Portfolio owner.

To enable access to the Portfolio data, the owner stores the encrypted transaction dataset, according to the OXT searchable encryption scheme on an external cloud server, and provides tokens to interested companies to perform searches on the encrypted Portfolio.

6.4.2 Assumptions and Setup

In this section we describe the setup functions needed for an individual to setup and maintain a Portfolio. We make the following assumptions:

- Each retailer has a certificate of their public key for signing the performed transactions (as described in the protocols presented in the following section).
- A reliable Internet connection is available to support communications during the Portfolio operations.

Credentials setup: An individual obtains credentials from the identity provider, according to the anonymous credential system described in Section 6.3.1.

Local Portfolio setup: The plaintext database DB is created and stored in the local Portfolio. According to the OXT scheme Setup function, a key K is chosen to be used for the keyword key and the identifier encryption. For each identifier, a pseudorandom label is computed, the identifier is encrypted and the label/ciphertext pair is inserted to a list L. After that, a dictionary D is built from L to operate as the server's index. To support dynamic additions to the dataset, an auxiliary encrypted database is created EDB^+, along with a dictionaryD^+ to which a pair (ℓ, d) is added with each keyword addition, where ℓ is a label computed from the keyword and a keyword-specific counter, and d is the encrypted record id added. Additionally, the dictionary D_{count} is maintained, associating each added keyword in the EDB with its counter value.

Cloud Portfolio setup: The individual registers with the external cloud service (E) using a pseudonym created from their credentials. The initial encrypted database EDB and dictionary D are uploaded, along with the empty auxiliary encrypted database EDB^+ and dictionary D^+ to support additions, according to the searchable encryption (OXT) scheme.

Retailer registration: The Portfolio owner registers with a retailer using a pseudonym created from the anonymous credential scheme (different and unlinkable to the one used for the cloud service). Transactions performed with the retailer will be introduced into the Portfolio.

Regarding the process to gain controlled access to Portfolio data, we consider the initial contact establishment between a Portfolio owner and an interested data consumer, whether a retailer interested to make an offer to the Portfolio owner or a data processing company, outside the scope of this work. Individuals could offer their Portfolio agent information directly to companies, to allow them to contact them with requests for data retrieval and related offers, or an external trusted service could handle the matching of Portfolios with certain verified characteristics to interested data consumers.

6.4.3 Portfolio Protocols

Two main protocols were designed to complement interactions in the Portfolio system. The first is the "transaction insertion protocol", that defines the way new transaction records are added to the individual's Portfolio when new transactions are conducted. The second is the "data retrieval protocol", that describes the data

retrieval process when a company wishes to access information from a Portfolio. Both protocols achieve the desired functionality, while offering privacy protection and information accountability. The protocols are discussed and analyzed in the following paragraphs.

Transaction Insertion Protocol

A new transaction is inserted into an individual's Portfolio according to the following protocol (illustrated in Figure 6.2):

1. When an individual makes a new transaction (T_i) with a shop, the shop signs it, creating $Sg_S(T_i)$, where Sg_S denotes the signature of the shop. The shop sends it to the customer's Portfolio agent, along with the relevant keywords for the transaction.

2. The individual stores the received transaction in the local Portfolio, creates the encrypted transaction and computes the labels for the transaction to be sent to the cloud service.

3. The cloud service receives the new encrypted transaction record, with identifier id, containing keywords W_{id} and their labels, and adds the received data into the user's auxiliary encrypted database EDB^+ and the corresponding auxiliary dictionary D^+.

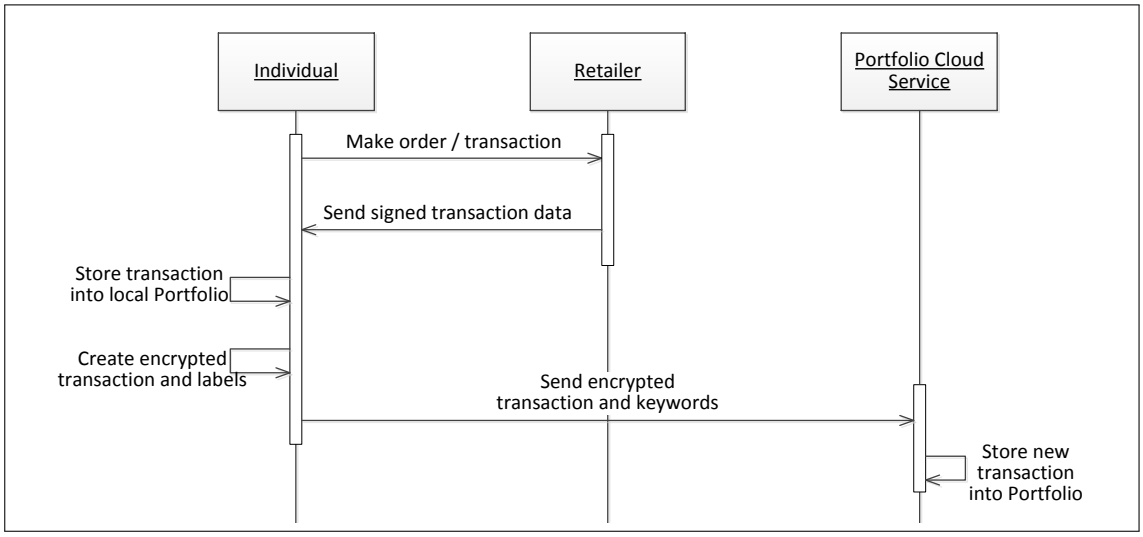

Figure 6.2: Basic transaction insertion protocol

Protocol additions

The above basic protocol ensures that transactions inserted in the Portfolio are

valid transactions and their contents have not been tampered with. To additionally provide assurances that the inserted transactions actually belong to the Portfolio owner, we propose the following (optional) additional steps to the insertion protocol:

- In Step 1 of the basic data insertion protocol, the retailer also sends the signed transaction hash value to the customer, to be used by the customer to prove the transaction ownership to the Portfolio cloud service.
- Additionally, the retailer also sends to the cloud service the signed transaction hash value and the user pseudonym used to perform the transaction with the retailer. This data will be used in the next step by the Portfolio cloud server to verify that the submitted encrypted transaction to be added into the Portfolio was performed by the Portfolio owner.
- In Step 3, the user sends, along with the encrypted transaction and its relevant keywords, the signed transaction hash value and a proof that the pseudonym used with the shop belongs to the same identity as the pseudonym used for the Portfolio service. This information is used by the cloud service to verify that the transaction belongs to the Portfolio owner, but is not inserted into the Portfolio data and will not be revealed during data retrieval.

Data Retrieval Protocol

The following protocol describes the way in which a company retrieves transaction data concerning specific keywords from an individual's Portfolio:

1. The company (C) contacts the Portfolio agent, stating the keywords they are interested in.
2. The Portfolio agent provides C a set of (homomorphically signed) query-specific trapdoors which the company can then transform into (homomorphically signed) search tokens as required by the OXT scheme. Note that the set of trapdoors given to C is fully determined by the query and independent of the searched data.
3. Using the query-specific search tokens, the company performs the relevant searches with the Portfolio cloud service. The result, produced using the OXT searchable encryption scheme, is the set of identifiers of the encrypted records matching the query, along with the associated record decryption keys.

4. The company requests the encrypted transactions from the cloud service and proceeds to decrypt them to access the transaction information.

Upon receipt of the authorised transaction information, the company can process it and offer personalised offers to the Portfolio owner.

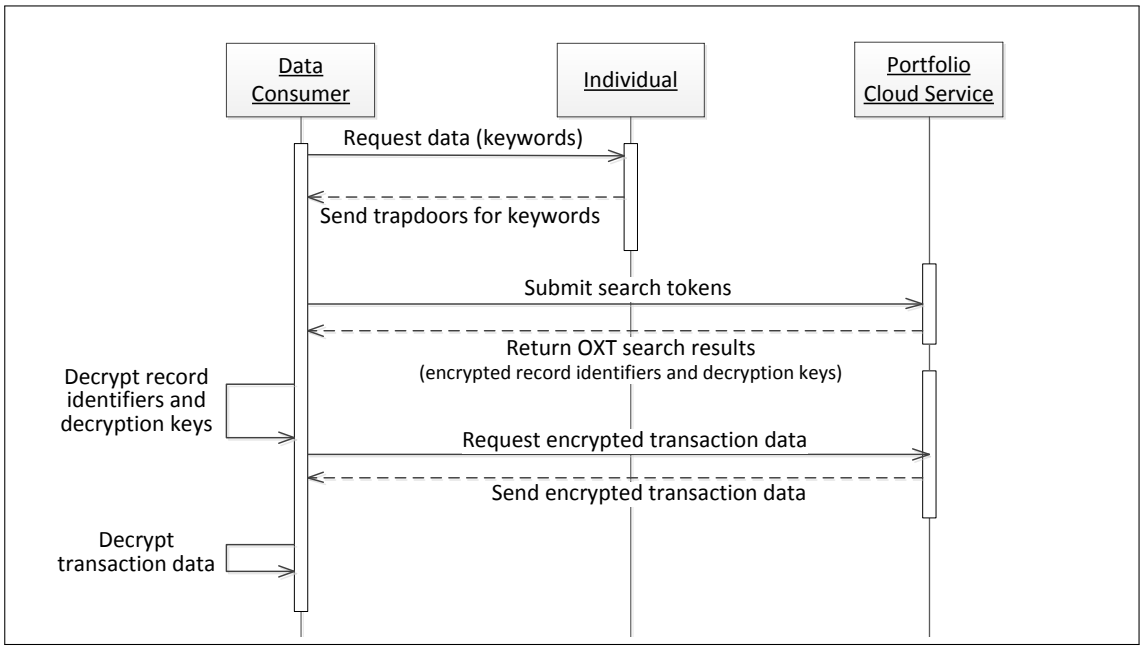

Figure 6.3: Data retrieval protocol

6.4.4 Security Discussion

The proposed architecture and protocols assume a semi-trusted, honest-but-curious, external storage server and arbitrarily malicious search-clients possibly colluding with each other, but not with the storage server. The security of the solution can be derived from the security properties of its components, namely the anonymous credentials and the OXT searchable encryption scheme.

The additional steps introduced in the data insertion protocol (Section 6.4.3) allow the honest-but-curious cloud server to keep a separate record of pseudonyms and signed transaction hash values in order to create a profile of the user, containing the shops the user purchased from and the pseudonyms used in those transactions. We consider this leak acceptable, given the accountability benefits these steps render to the solution.

Additionally, due to the use of the OXT searchable encryption scheme, there exists leakage of data access and query patterns to the storage server, but no direct

plaintext data or search values are leaked. More specifically, information on the number of records matching one of the query terms (typically the least frequent one) is leaked. This leakage is inevitable in order to maintain the searchable encryption efficiency [35] and may be improved with future searchable encryption scheme developments.

6.5 Discussion

In this work we propose an alternative approach to the way consumer data is handled by moving the control of the transaction profile back to the consumer. The availability of applicable cryptographic techniques allows for the current practice on consumer profile management to improve, offering capabilities for increased data protection and opportunities for trustworthy utilization of consumer profiles, applicable in the existing common infrastructure. The main contribution of this work is the construction of an integrated, practical solution that illustrates the applicability of the approach. The proposed architecture utilizes contemporary cryptographic building blocks which are combined within the developed protocols to achieve the intended functionality.

We expect that this new paradigm will appeal to individuals, who will be able to retain control of and utilize their personal information, as well as well-meaning retailers who will enhance their company profile by adopting this privacy-sensitive practice and additionally, will be relieved of the burden of maintaining and protecting customer data. Moreover, we expect this solution to appeal to data consumers interested in accessing accountable, up-to-date consumer profiles, without the risk of violating the consumer's privacy. Even retailers that are not interested in maintaining and processing consumer profile data will have incentive to participate, as Portfolio owners will prefer making transactions with them, to enrich their consumer profiles.

Our approach allows for retailers and service providers to retain anonymous information on transactions individuals conduct with them, but also adds the opportunity to gain authorised access to accountable information on the customer's transactions with other retailers. On the other hand, individuals retain control of their consumer profile and are able to provide controlled access to it, whenever they choose. Motives for doing so include personalised special offers and privileges, or even integration with a personal data market, incorporating a compensation system such as FPIT (Chapter 4).

Chapter 7

Privacy-preserving, User-centric VoIP CAPTCHA Challenges: an Integrated Solution in the SIP Environment

7.1 Introduction

The audio CAPTCHA challenges used today to prevent automated Spam Over Internet Telephony (SPIT) attacks on Voice over IP (VoIP) services do not take into account the characteristics of the caller or the callee. The fact that these challenges are generic, does not allow for the process to take into consideration the cognitive abilities of human users, while at the same time discriminates against users that have difficulties solving the generic challenges. The ability to use information about the caller or the callee opens up new opportunities for creating more effective and fair CAPTCHA challenges. However, the required information about the caller will probably be sensitive personal information, and thus it is important that a privacy-preserving method of achieving the adaptation of CAPTCHA challenges is used.

In this work we argue that it is possible to address discrimination issues that naturally arise in contemporary audio CAPTCHA challenges and potentially enhance the effectiveness of audio CAPTCHA systems by adapting the challenges to the user characteristics. We design the PrivCAPTCHA prototype to offer privacy-preserving, user-centric CAPTCHA challenges. Anonymous credential proofs are integrated into the SIP protocol and the approach is evaluated in a real-world VoIP environment. The results of this work indicate that it is possible to create VoIP CAPTCHA services offering privacy-preserving, user-centric

challenges, while maintaining sufficient efficiency. The proposed approach was evaluated through an experimental implementation to demonstrate its feasibility.

PrivCAPTCHA aims to achieve fairer, non-discriminating CAPTCHA services, while protecting the user's privacy. Adoption success relies upon the general need for employment of privacy-preserving practices in electronic interactions. The purpose of this approach is to enhance the quality of life of users, who will be able to receive CAPTCHA challenges closer to their characteristics. This applies especially to users with disabilities. Additionally, as a privacy-preserving service, this approach is expected to increase trust during the use of services that employ it. To our knowledge this is the first comprehensive proposal for privacy-preserving CAPTCHA challenge adaptation. The proposed system aims at providing an improved CAPTCHA service that is more appropriate for and trusted by human users.

7.2 Related work

CAPTCHA challenge solving is a highly user-dependent process. Works on the subject illustrate that the user characteristics can greatly affect the success rate of the mechanism usage. As described in the following paragraphs, the user's sensory and cognitive abilities, computer literacy and language fluency play a decisive role on the user's ability to solve CAPTCHA challenges. In this section, we briefly survey the audio CAPTCHA technology, which is in the early stages of development. We focus first on evaluating the existing audio CAPTCHAs and then on recording the CAPTCHA attributes which improve its usability.

7.2.1 Existing CAPTCHA Evaluation

Existing audio CAPTCHAs have been proven more difficult to use for visually impaired than non-visually impaired people [20]. For their research they used 162 persons, of whom 89 were visually impaired, and popular website audio CAPTCHA implementations. Their research illustrated that audio CAPTCHAs are difficult to solve. Only 43% of users with visual impairments were able to answer an audio CAPTCHA at the first attempt and only 39% of other users. Moreover, it should be noted that visually impaired users took at least twice as long. Yet nearly half of the users (47%) still failed to respond correctly to an audio CAPTCHA after 3 attempts. This is a somewhat unexpected result, since one would anticipate that audio CAPTCHA challenges would be more appropriate for visually impaired persons.

Bursztein et al. [24] conducted an extensive study on the ability of people to

solve existing CAPTCHAs, as well. Regarding audio CAPTCHAs, they studied eight of the most popular implementations. The conclusions that emerged from their study were a) the period for listening and solving a CAPTCHA is certainly excessive (averaged over 25 seconds), b) the percentage of users who took second or third attempts, because the previous attempt was wrong, exceeded 50%, and c) people who were not native English speakers had major problems in solving the CAPTCHA and therefore the success rate was reduced by more than 20%.

Soupionis and Gritzalis [130] classified the audio CAPTCHA attributes, evaluated the current popular audio CAPTCHA implementations and developed a new audio CAPTCHA for VoIP environments. The CAPTCHAs were classified based on their attributes into four categories: (a) vocabulary, (b) background noise, (c) time, and (d) audio production. Afterwards, the evaluation took place where the CAPTCHAs were utilized on the above mentioned attributes. The evaluation process was based on the fact that CAPTCHAs must be easy for human users to solve, easy for a tester machine to generate and grade, and hard for a software bot to solve. Therefore, the final evaluation was made by two means; namely, by user tests (60 persons) and by two bots configured to solve audio CAPTCHAs. The evaluation process proved that a) the current CAPTCHA implementations are not adequate, meaning that every implementation is either too easy or too difficult to be solved by both users and bots, and b) the implementation attributes of some CAPTCHAs, like long vocabulary ($>$ 8 characters) and language requirements (native vs. non-native English speakers), negatively affects the users' success rate (40%) in most cases.

7.2.2 CAPTCHA Usability

Yan and El Ahmad [146] discuss usability issues that should be considered and addressed in the design of CAPTCHAs. The authors analyze a few aspects for the three main types of CAPTCHAs: 1) Text-based, 2) Sound-based, and 3) Image/picture-based schemes. As far as the Sound-based ones are concerned, the major parameters contributing to having a usable CAPTCHA are the following:

1. Distortion, meaning the background noise which distorts sounds in audio CAPTCHAs.

2. Content, meaning the content materials used in audio CAPTCHAs which are typically language specific, like digits, character set and string length.

3. Presentation, meaning the integration technique within the web pages which is still a great concern.

Lazar et al. [97] describe the development of a new audio CAPTCHA called the SoundsRight CAPTCHA focusing on blind users. The authors identify that

one of the major problems is the "linear audio CAPTCHA playback" problem. A blind user has to quickly navigate through a page using only keyboard strokes, and the screen reader causes audio interference with the audio CAPTCHA that is being played. Once the audio CAPTCHA is played, the user must quickly focus on the answer portion of the CAPTCHA to input what they heard. This problem may affect mainly the web implementation of audio CAPTCHA, but it shows that there should be a certain method for the CAPTCHA playback. The above results indicate that there is a need and the potential to create more appropriate challenges for the human user that will allow for fewer problems in solving CAPTCHA challenges. Additionally, user-centric challenges provide the potential to weaken the connection between the difficulty of CAPTCHA solving for humans and for bots respectively, resulting in more effective CAPTCHA systems.

Taking into account the person's characteristics during CAPTCHA generation, brings about privacy concerns that need to be addressed. There has been significant progress on the subject of accountable, privacy-preserving services during the past decade. The privacy-preserving techniques used in the proposed system are closely related to accountable anonymous communication systems [49], anonymous credential systems [33] and electronic identity cards [114, 48]. Using cryptographic tools, all these systems aim at providing their functionality while protecting users' privacy. Similarly, we utilize existing cryptographic primitives to create a privacy-preserving personalised CAPTCHA system, which allows users to prove attributes about themselves and receive personalised challenges, without revealing their identity.

7.3 Concepts and Motivation

The selection of an appropriate CAPTCHA challenge that successfully distinguishes between human users and bots is a challenging task. Generic CAPTCHA challenges that are difficult enough for bots, often pose difficulties to human users as well. Therefore, a method is needed to tailor CAPTCHA challenges closer to the human user, without necessarily lowering the difficulty level for bots. Overall, this work does not aim at making CAPTCHA challenges generally easier, rather, it aims at proposing a method to create more appropriate, effective and fair challenges for the users. However, the process of selecting the appropriate kinds of adaptations according to the user's characteristics is outside the scope of this work. We propose a privacy-preserving method of proving user characteristics to the CAPTCHA service and delivering the adapted challenge. We demonstrate the feasibility of the proposed system through representative examples.

One could argue that proving user characteristics to the VoIP service eliminates the need for the CAPTCHA test overall. However, we believe that the

CAPTCHA test is still needed to protect users from unauthorised use of their accounts (hijacking) and attempts to impersonate them. Additionally, the combination of anonymous credential proofs and the CAPTCHA test, protects the VoIP system from the unauthorised use of credentials and from users that misuse their credentials for making SPIT calls. Therefore, this work does not aim at making the CAPTCHA test obsolete through the use of anonymous credential proofs. Using the PrivCAPTCHA approach, callers can assert that they are human users and have certain characteristics, but this is also verified during the CAPTCHA test.

As with many IT applications and services, CAPTCHAs were initially implemented for the average or typical user, without catering for users with special needs and (dis)abilities. A user may not solve a CAPTCHA test due to their personal characteristics, e.g. their language, age, mental or physical disabilities. This person faces discrimination and violation of their communication rights. It is important to ensure that services are accessible by the majority of the population, regardless of their abilities and characteristics. Discrimination issues concerning CAPTCHA challenges are discussed in more depth in [133]. Taking into account user characteristics to personalise the VoIP CAPTCHA service, addresses effectively the discrimination concerns. This solution, however, introduces new privacy requirements, as the personal information used to provide the personalised CAPTCHA service, need to be protected from unauthorised access. Individuals' privacy rights determine that they should be in control of their information, being able to determine when and who acquires knowledge about their preferences and characteristics.

7.4 PrivCAPTCHA Architecture

In this section the components of the proposed architecture are presented first, comprising the cryptographic building blocks used to achieve the privacy properties and the entities that participate in the system. Then, the resulting functionality is described.

7.4.1 Cryptographic Building Blocks

In order to achieve the privacy-preserving attributes of PrivCAPTCHA, we use anonymous credentials and a data management unit as cryptographic building blocks. In this section we provide a high level description of these building blocks, focusing on their attributes and functionality.

Anonymous credentials

Anonymous credentials [26] allow users to acquire credentials and demonstrate them without revealing their identity. Using the private credential system described in [33], individuals can use different unlinkable pseudonyms, based on the same credential issued by an identity provider. The private credential system can also provide certified attributes by the identity provider, for the individual to selectively reveal attributes (e.g. their age range, based on their date of birth). Anonymous credentials constitute today an accepted and applicable privacy enhancing technology. In this work we use Idemix[1], an open source anonymous credential system (see implementation in Section 7.5.1).

The entities participating in an anonymous credential system are individuals, companies and trusted third parties (e.g. government services), which can assume the roles of issuers, recipients, provers and verifiers. A credential is created by an issuer for a recipient by executing the issuing protocol. The recipient (i.e. the credential owner) can then create a credential proof, to be used by a verifier to verify the validity of the credential (proving protocol).

To be able to issue anonymous credentials, an issuer needs to generate a public key pair and create specifications of the structures of the credentials issued. These specifications and the issuer's public key are then published to be used during the proof protocol.

In order to acquire an anonymous credential, a user chooses a master secret key, according to the agreed upon system parameters (bit length, groups to be used). This secret key enables the creation of multiple unlinkable pseudonyms by the user, to be used with different service providers (in our case VoIP services). The issued credential consists of the issuer's public key, the credential structure (necessary for the verification of its validity) and the attribute values.

During the proving protocol, the prover (i.e. the credential owner) creates a proof on behalf of the verifier that proves ownership of a certain credential. The verifier checks the validity of the given proof. Credential attribute values contained in the proof may or may not be revealed to the prover, according to the settings of the proof creation process.

A local data storage and management unit

A data management unit that resides at the owner's side, similar to the Portfolio architecture proposed in [132] is used to manage the user's credentials and certificates. The contents of a user's portfolio include:

- Anonymous credentials, containing verified demographic data and personal characteristics, e.g. age, education level, disabilities.

[1]Identity Mixer, http://www.zurich.ibm.com/idemix/usage.html

- Certificates of successful CAPTCHA tests issued by the CAPTCHA service. This transaction history can be used to provide further evidence to the CAPTCHA server that the user is human and non-malicious and can even be used to allow users to pass over the CAPTCHA test for a limited time.

7.4.2 Entities in PrivCAPTCHA

The entities that participate in the proposed system are the following:

The Identity Provider (IDP). Users obtain their credentials from the IDP, by registering an identifier (e.g. their social security number) and a pseudonym P. The IDP is considered a trusted third party (like a passport authority) that retains the user information together with their pseudonym. The IDP does not need to be a single entity, it can be a distributed service to achieve better service availability and enhanced security.

The User. In our system the users are considered the VoIP service users. All can act both as callers and callees. When acting as callers, their portfolio information can be used to receive personalised CAPTCHA challenges as illustrated in Figure 7.1.

The CAPTCHA server, which acts as a verifier for the anonymous credential system. Moreover, the CAPTCHA server automatically generates the CAPTCHA challenge, according to the proven user attributes and evaluates the provided answer.

The entities of the PrivCAPTCHA system and their interactions are shown in Figure 7.1. The user's profile is stored at the personal portfolio residing at the user's side. After registering with an IDP and obtaining the anonymous credentials, the caller can prove some attributes to the CAPTCHA server and receive personalised CAPTCHA challenges. According to the proportionality principle the IDP retains no more data than that strictly required to serve the personalised CAPTCHA service. The IDP entity combines the retained data with a pseudonym in order to protect the identity of the user and it is not allowed to reveal or to use this data for any other purpose, with the exception of law enforcement purposes if required and to the extent that is provided by the respective law. The proposed architecture provides personalised and effective VoIP CAPTCHAs while preserving the privacy and communication rights of the user.

7.4.3 System Functionality

In this section we will describe the general functionalities of the system. Regarding the cryptographic primitives used (see Section 7.4.1), we adhere to their descriptions as proposed by their authors and we only provide descriptions of their use within the context of our work.

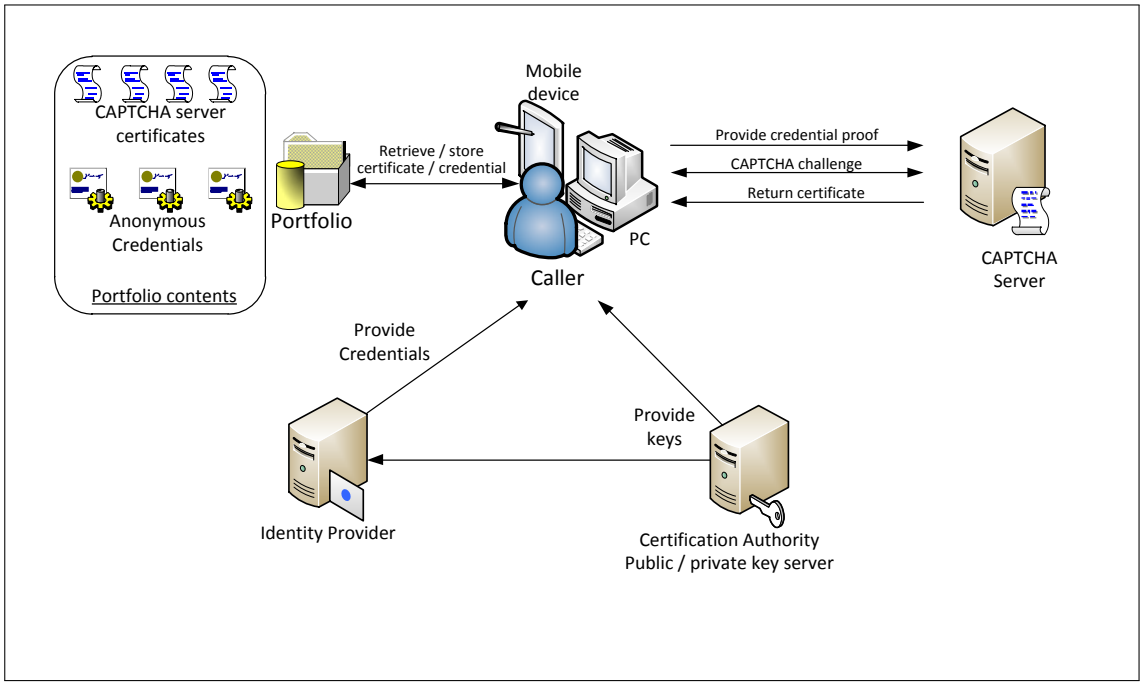

Figure 7.1: The PrivCAPTCHA architecture

Proving attributes to the CAPTCHA system. After acquiring the anonymous credentials by the IDP the user can begin using them to prove attributes to CAPTCHA services in order to receive personalised challenges. The user needs to prove to the CAPTCHA service that the credential is valid, which verifies that the user has a certain attribute. The verification mechanism is based on efficient zero knowledge proofs [33]. In the context of VoIP communications, the credential proof can be incorporated into the SIP INVITE message and transmitted seamlessly during the VoIP call (see details in Section 7.5.3).

CAPTCHA generation and outcome. After receiving and verifying the user's credential proof, the CAPTCHA service generates the appropriate CAPTCHA challenge, according to the user characteristics. The characteristics that the CAPTCHA server takes into account in our example implementation are:

Disability status. Depending on the user's disability, the system can adapt the CAPTCHA challenge to provide a more appropriate challenge for the user's abilities. For example, hearing impaired users may need an elevated volume level for the challenge.

Language requirements. Based on the user's native language, the CAPTCHA server can provide the appropriate challenge, either a challenge in the user's

language, or some adaptation to facilitate non-native speakers of the challenge language.

Age. The age of the user affects the ability to solve tests. If the user is too young or too old, then the CAPTCHA challenge should be adapted accordingly, e.g. to contain less characters to be recognised or to allow more time to solve the test.

Apparently, additional user characteristics can be considered apart from the aforementioned ones, like education level (e.g. literacy), learning disabilities (e.g. dyslexia), etc. The selection of appropriate adaptations for the challenge to facilitate each user characteristic is outside the scope of this work. Our goal is to illustrate the mechanism for the adaptation according to any given user characteristic.

CAPTCHA server certificates. Upon successful completion of a CAPTCHA challenge, the CAPTCHA server sends the user a certificate, attesting that this user did make a legitimate communication. This certificate is stored into the user's portfolio and can be used later by the same user to prove previous legitimate use of the CAPTCHA service.

The above functionalities are prone to misuse and malicious behavior on the part of the user. In Section 7.4.4 we address the main issues that have been identified for the proposed system. The communications protocol in PrivCAPTCHA contains the following steps:

1. The user makes a SIP call, containing the credential proof (see Section 7.5.3).
2. The CAPTCHA server verifies the provided credential proof.
3. The CAPTCHA server generates the appropriate challenge.
4. The user responds to the CAPTCHA challenge.
5. If the challenge response is correct, the CAPTCHA server sends the user a certificate of successful completion.

7.4.4 Incident Response Requirements

Designing a system on a user-centric driven security basis requires that the system is robust in the sense that the user is not significantly exposed during a security failure. In the proposed system we have identified the following aspects and requirements for incident response and escalation procedures in the event of a security failure.

Tolerance to false positives. We can in principle consider that false positives carry a minor security impact. The event of a bot successfully answering the audio CAPTCHA challenge will be detected by the destination/callee and the service should maintain the facility for the callee to report/redirect the call for further logging and analysis. Responding to false positives is a good example of active user participation in the security process. Regarding CAPTCHA server certificates, in case of false positives, a revocation procedure can be followed upon receipt of the callee report.

Tolerance to false negatives. Rejecting a legitimate call request after a failed audio CAPTCHA attempt is an event of major significance. Therefore the underlying security parameters are expected to be set on a level where the false negatives are minimised despite the drop in security. Indeed, giving priority to user acceptance over security is part of the user-centric system design practice. In addition, there needs to be a continuous evaluation similar to vulnerability assessment practices. More specifically, as a security administrator must be informed and proactively search for new vulnerabilities affecting the system, the audio CAPTCHA engineer must keep the system up to date with the state of the art research in order to maintain the optimum level of security versus user acceptance.

Tolerance to CAPTCHA server certificate misuse. The certificates provided by the CAPTCHA server can be (un)intentionally misused by users to exploit the system. In case of reported malicious use of these certificates, revocation methods [26] can be examined.

Correlate system failures with SPIT results. A threat management system should be implemented and the audio CAPTCHA service should be placed in the wider system security context in order to identify threat vectors that may target the CAPTCHA but also exploit the system as a whole.

Reputation management. Reputation mechanisms introduce a number of security issues and should these become part of the audio CAPTCHA service, reputation misuse should be addressed with well defined escalation procedures. The proposed system can adopt published procedures and controls for reputation management.

7.5 Implementation and Tests

To demonstrate the feasibility of our proposed system, we created an implementation that integrates the core PrivCAPTCHA functionality into a real-world, open source VoIP application. This implementation allows a credential proof to be transmitted to the CAPTCHA server within a SIP call. For this purpose we created a SIP custom header containing the credential proof for the CAPTCHA

server to receive and verify and introduced it into the SIP INVITE message. Additionally, to allow the CAPTCHA server to verify the credential proof, we created a ProofVerifier executable that is called when the SIP message with the custom header is received. In the following sections we describe the developed implementation and the experimental results from its execution.

7.5.1 The PrivCAPTCHA Anonymous Credentials

The *Identity Mixer* cryptographic library[1] was used to create the anonymous credentials for the CAPTCHA user and the proof to be sent to the CAPTCHA server. The library implements the anonymous credentials of Camenisch and Lysyanskaya [31], i.e. the functionality of anonymous authentication for the issuer, client, and service provider.

To create the PrivCAPTCHA anonymous credentials for the implementation we used the Idemix library (Release 2.3.4). We created the appropriate credential structure and proof specification for our application and, utilizing the library functionality, added implementations for the issuance, proof creation and verification methods of the PrivCAPTCHA credentials.

As mentioned in Section 7.4.3 , the credentials created for this implementation describe the following credential-owner characteristics and corresponding enumerated attributes:

- Disability status: Hearing Impaired, Blind, Illiterate
- Native language: English, other
- Age group: Child, Adolescent, Adult, Elderly

Further categories and possible values can be added to describe user characteristics, to suit the needs of each application.

The credential structure used in our implementation is presented in Figure 7.2, following the credential annotation described in [19]. The credential information is partitioned into the attributes, defined by a name, issuance mode and type of attribute, and the implementation, where implementation specific information is provided. The enumerated attributes are implemented by assigning a distinct prime to each possible attribute value, according to [28].

In Figure 7.3 we present an example credential in accordance to the PrivCAPTCHA credential structure (Figure 7.2):

To implement the proving protocol, we needed to create a proof specification, shown in Figure 7.4, for our example credential (Figure 7.3).

[1] https://prime.inf.tu-dresden.de/idemix/

```
Attributes{
Attribute { Status, known, type:enum }
   { HearingImpaired, Blind, Illiterate }

Attribute { NativeLanguage, known, type:enum }
   { English, Other }

Attribute { AgeGroup, known, type:enum }
   { Child, Adolescent, Adult, Elderly }

Implementation{
PrimeFactor { Status: HearingImpaired = 3 }
PrimeFactor { Status: Blind = 5 }
PrimeFactor { Status: Illiterate = 7 }
PrimeFactor { NativeLanguage: English = 11 }
PrimeFactor { NativeLanguage: Other = 13 }
PrimeFactor { AgeGroup: Child = 17 }
PrimeFactor { AgeGroup: Adolescent = 19 }
PrimeFactor { AgeGroup: Adult = 23 }
PrimeFactor { AgeGroup: Elderly = 29 }

AttributeOrder { Status, NativeLanguage, AgeGroup}
}
```

Figure 7.2: PrivCAPTCHA credential structure

```
References{
Schema=http://privCAPTCHAdomain.com/credCAPTCHA.xsd
Structure=http://privCAPTCHAdomain.com/CredStructCAPTCHA.xml
IssuerPublicKey=http:// privCAPTCHAdomain.com/exampleIPK.xml
}

Elements{
Signature { A:..., v:..., e:... }
Values { Status:HearingImpaired; NativeLanguage:English; AgeGroup: Elderly }
}
```

Figure 7.3: PrivCAPTCHA example credential

```
Declaration{ id1:revealed:enum; id2:revealed:enum;
             id3:revealed:enum;}
ProvenStatements{
   Credentials{
     randName1: http://privCAPTCHAdomain.com/CredStructCAPTCHA.xml=
        { Status:id1, NativeLanguage:id2, AgeGroup:id3 }
     }
}
```

Figure 7.4: PrivCAPTCHA proof specification

The proof specification contains:

- Credentials that the user proves ownership of.
- Identifiers for the values included in the proof. Identifiers are assigned to attributes that are fully or partially revealed.
- Attribute types of each identifier, which need to match during the proof protocol.
- Constants that can be assigned to identifiers.

7.5.2 Jitsi - Open Source VoIP Application

To integrate the SIP custom header containing the credential proof into the VoIP call, we used Jitsi[1], an open source multi-platform audio/video Internet phone application. It supports several instant messaging and telephony protocols, including the Session Initiation Protocol (SIP) used in VoIP networks. Jitsi and its source code are released under the terms of the LGPL. Jitsi is mostly written in Java and, among others, it uses the JAIN-SIP protocol stack for SIP support.

For the purpose of our implementation we downloaded the Jitsi v2.2.4603.9615 source snapshot and added code to introduce the custom header containing the credential proof into the SIP INVITE message.

7.5.3 SIP Custom Header

For the credential proof to reach the CAPTCHA server, we introduced the CredentialProof custom header into the SIP INVITE message sent by Jitsi during the SIP call. To achieve that, the class CredentialProof was created, extending

[1] https://jitsi.org/Main/Features

the ParametersHeader class from the JAIN-SDP library[1], which implements the parameters setting functionality of the SIP headers. Using the createRequest method of the SipMessageFactory class in the "net.java.sip.communicator.impl.protocol.sip" package, the custom header was appended into the SIP INVITE message header.

Using this implementation, each time a Jitsi user makes a call, the CredentialProof header is transmitted to the recipient (in our case the CAPTCHA server). The CredentialProof header is shown in Figure 7.5, where the SIP call performed by Jitsi is captured with the Wireshark packet analyzer[2].

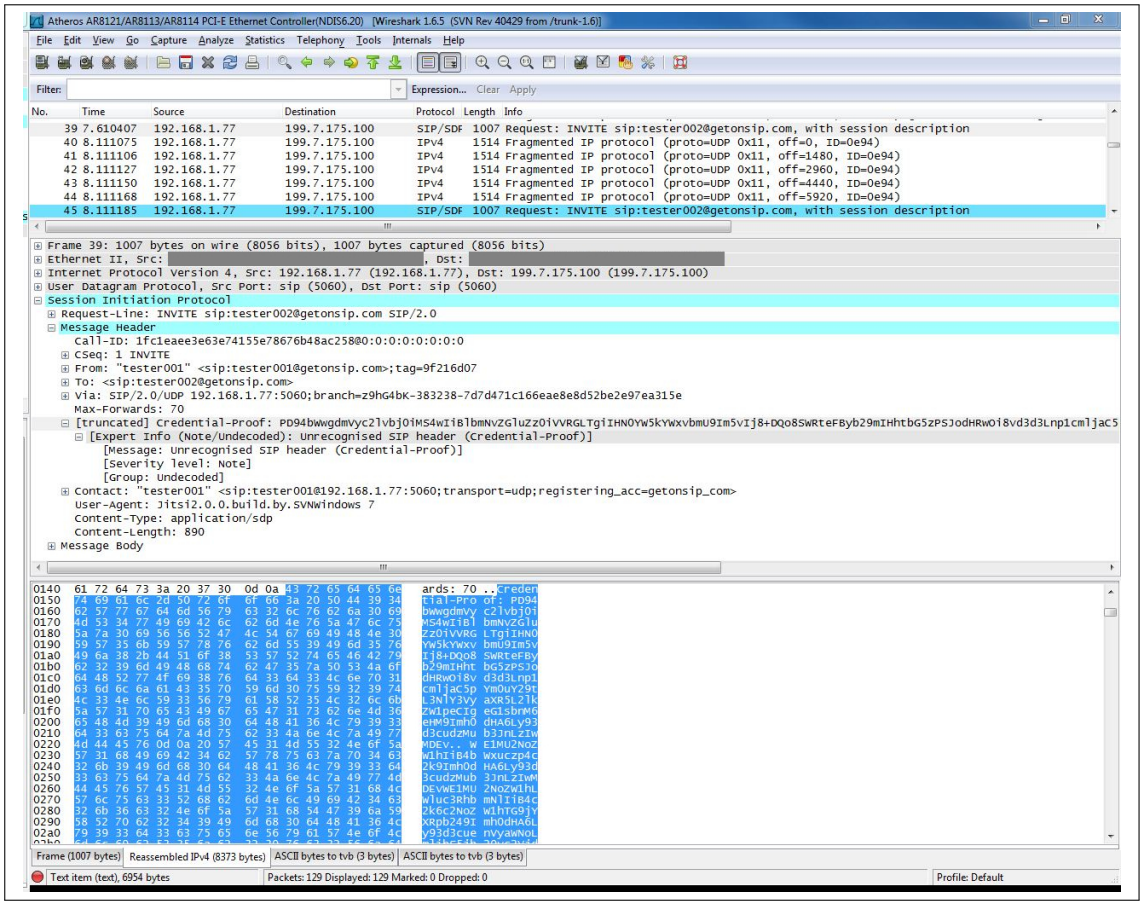

Figure 7.5: CredentialProof custom header captured with Wireshark during SIP call

[1]JSIP: Java API for SIP signalling, https://jsip.java.net

[2]http://www.wireshark.org/

7.5.4 Experimental Network Environment

The network environment (Figure 7.6) that we have implemented consists of the following fundamental entities:

An Asterisk server, which was based on AsteriskNow[1], a widespread open source SIP server implementation, also used for VoIP PBX. The provided API of the server supports easy manipulation of SIP headers and allows storing useful metadata in the call-records database. The VoIP PBX runtime environment offers administration access via command line or through the FreePBX web-based application over an Apache server and includes a MySQL database, that stores operational parameters, such as SIP extensions, voice trunks, call records, etc. The implemented server has been customised in order to register users, redirect SIP messages and establish calls. The PC used for setting up the Asterisk server was a Pentium 4, 2.8GHz with 2GB RAM.

A CAPTCHA service. The audio CAPTCHA was implemented as a separate service for efficiency reasons, since the computational resources needed for such a module are considerable. The service was implemented on the AsteriskNow software as well. The basic algorithm was developed using the PHP class Asterisk Gateway Interface (PHPAGI[2]), which interacts with the AsteriskNow software to provide audio CAPTCHAs as a standalone service. The CAPTCHA service:

1. receives the SIP message,
2. extracts the values of the custom header,
3. passes the values to the PHPAGI module,
4. identifies the characteristics of the user asked to solve the CAPTCHA,
5. selects and ”plays” the appropriate audio CAPTCHA file based on the proven characteristics,
6. validates the answer and either sends the decision to Asterisk server or it re-sends a new CAPTCHA.

Various VoIP callers. These callers are programmed to make calls to the VoIP service clients. In our scenarios, they are redirected through the CAPTCHA service. The exact number of the external callers depends on each use case/scenario.

[1] http://www.asterisk.org/downloads/asterisknow
[2] http://phpagi.sourceforge.net/

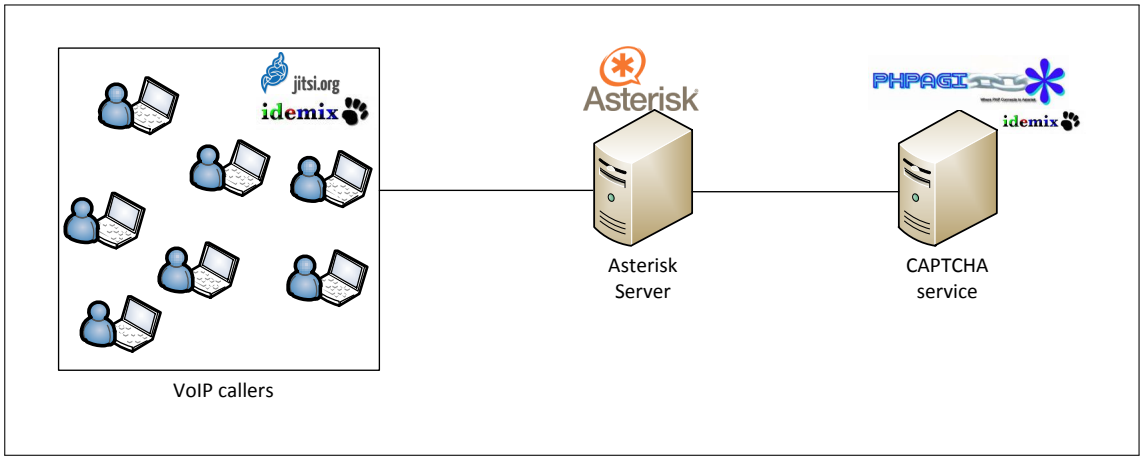

Figure 7.6: Experimental network environment

7.5.5 Proposed Scenarios and Results

For the experimental process we created two scenarios. In both scenarios we measured the time needed for the CAPTCHA server to select the appropriate audio file and set the right playback attributes (volume, number of retries, etc.).

The first scenario was to have a single call initiated and measure the aforementioned needed time. The time needed without the PrivCAPTCHA validation was 3.6ms and with it was 7.8ms. Even though the time has been doubled it still stays extremely low, so there is no significant overload to the system.

The second scenario consisted of 20 external clients, which initiated new calls. The SIP calls were generated randomly, but there was a limit of a maximum of 20 simultaneous calls. The average time interval between the calls was 100ms. The total number of calls was 460. Finally, it should be stated that the calls were terminated 10 seconds after the call establishment, while playing the CAPTCHA audio file. This means that the CAPTCHA service has to keep the calls established either for 10 sec or until the CAPTCHA timeout is reached, which is based on the chosen characteristics. This scenario was created using the SIPp[1] call generator.

In Figure 7.7 we present the results of the second scenario. The min and max values of each measured variable are represented by the top ($\top$) and bottom ($\bot$) bars. The mean value and one standard deviation from it are represented by the (-) bar and the greyed box respectively. The first column depicts the time need for the proof to be verified by the CAPTCHA server, i.e. the delay caused by the PrivCAPTCHA addition in the system. The second column represents the total time needed for the request to be processed and the third column the time needed for the CAPTCHA challenge to be generated.

[1] http://sipp.sourceforge.net/

The results show that the proof verification process requires approximately 300ms, with relatively narrow deviation, which we consider to be low enough to make this addition a feasible option. Additionally, the efficiency of this verification can be further improved via a server to amortise Java Virtual Machine startup costs, which in this experiment were calculated to be approximately 100ms.

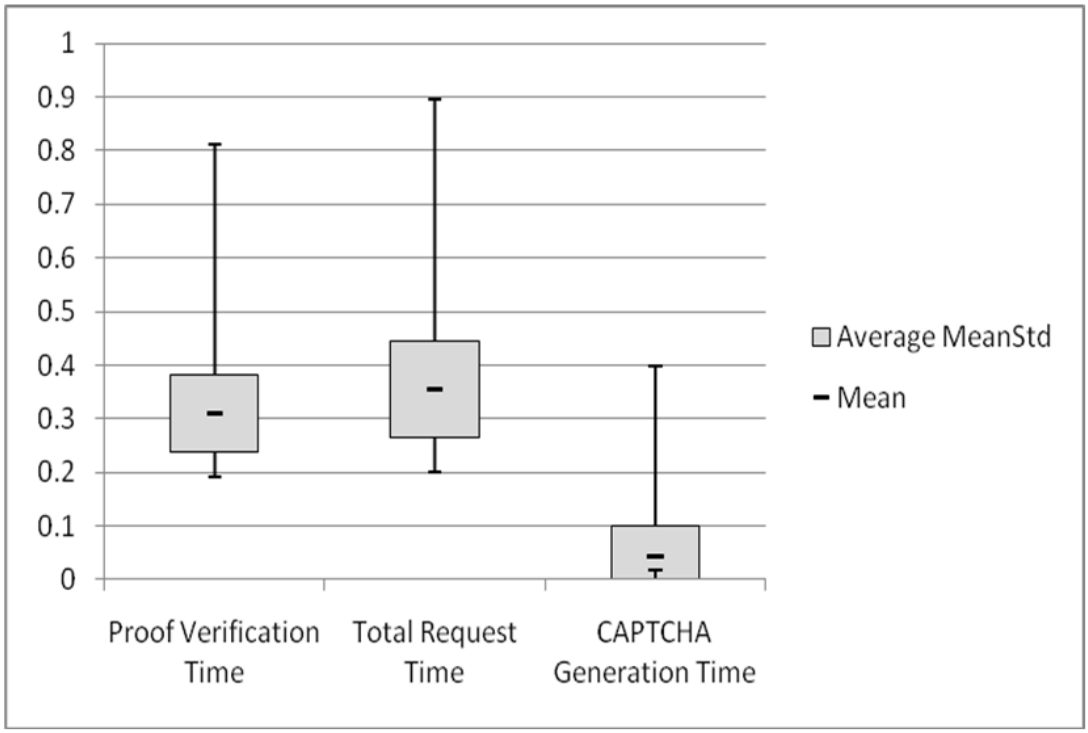

Figure 7.7: Experimental performance results (in seconds)

This experiment was conducted to demonstrate primarily the feasibility of the proposed system, not its efficiency. Therefore, the times recorded indicate an upper bound on the processing time. In case PrivCAPTCHA were to be used in a production CAPTCHA service, the efficiency of the system can be expected to be enhanced, by removing overheads and setting up more than one CAPTCHA servers.

Nevertheless, it is a fact that the addition of PrivCAPTCHA introduces an additional computationally intensive task to the SIP communication process, decreasing the CAPTCHA server's maximum capacity for concurrent connections. This can be exploited to enhance the effectiveness of attacks against the SIP protocol, an existing problem of VoIP communications [45, 90]. Since the additional header introduced is created according to the SIP protocol specification, the same countermeasures and policies used to protect conventional VoIP systems can be

employed to protect the PrivCAPTCHA system [129, 62, 57]. These can reduce the amount of malicious messages that reach the CAPTCHA server and invoke the header verification process.

For example, possible countermeasures include SIP message evaluation mechanisms, such as a SIP message parser, which will evaluate the CredentialProof header content. In case of an unsophisticated attack, during which the header is used without maintaining the valid xml schema of the proof, the message will be filtered by the SIP parsing mechanism, therefore never reaching the CAPTCHA server. During a more sophisticated attack, wherein the required xml schema of the proof is valid and only the proof components are not, or a valid CredentialProof is replayed, the proof validation mechanism will be invoked, increasing the server's computational load and possibly leading to service downtime. In this case logging and blacklisting mechanisms can be used to block the malicious calls and server resource monitoring can be used to disable the CAPTCHA service and prevent incontrollable VoIP service downtime when resources are depleted until the attackers are blacklisted.

7.6 Discussion

In this work, we propose a user-centric, privacy-preserving VoIP CAPTCHA adaptation approach. The PrivCAPTCHA architecture combines existing cryptographic technologies, which provide strong privacy guarantees, utilized under a new context. The proposed system aims at providing an improved CAPTCHA service that is more appropriate for and fair to the human users. Descriptions of the proposed system functionalities are provided and an experimental implementation is carried out within the VoIP protocol. Experimental results show that the utilization of cryptographic tools introduces an additional computational overhead into the audio CAPTCHA application. We consider this overhead to be tolerable for modern computational platforms possibly combined with appropriate performance optimisation techniques. However, this overhead can be exploited to increase the efficiency of attacks against the SIP protocol, therefore appropriate countermeasures should be employed to decrease these negative effects. Moreover, although we mainly consider CAPTCHA challenges for VoIP calls in this work, we believe that this idea can be useful for providing a general mechanism for CAPTCHA adaptation according to the user's characteristics.

Chapter 8

myVisitPlanner: Privacy-preserving Personalised Itinerary Planning System for Tourism

8.1 Introduction

MyVisitPlanner is an intelligent web-based application aiming at making recommendations that help visitors and residents of the region of Northern Greece to plan their activities during their stay in this area [120]. The system encompasses a rich ontology of activities, categorised across dimensions such as activity type, historical era, user profile and age group. Each activity is characterised by attributes describing its location, cost, availability and duration range. The system makes activity recommendations based on user-selected criteria, such as visit duration and timing, geographical areas of interest and visit profiling. A recommendation engine employs non-intrusive machine learning techniques to dynamically update user profiles, concerning user preferences for both activities and resulting plans, while taking privacy concerns into account.

Personalised services, such as myVisitPlanner, by definition need user characteristics and preferences to improve their service quality. The more data the system has on the user, the more effective personalisation of the service can be achieved, closer to the user needs. On the other hand, user privacy protection requirements contradict the availability of unrestricted access to detailed user data. Achieving a balance between sufficient privacy protection and efficient personalisation is the key for a successful privacy-preserving personalisation system.

Sevice personalisation in myVisitPlanner is achieved by utilizing information

regarding user preferences and characteristics, acquired directly from the user and indirectly by machine learning techniques. This information, although crucial for achieving successful recommendations, needs to be protected from unauthorised access and use. Privacy protection mechanisms, incorporated in the system design, enhance user trust toward the service and relieve system administrators from the responsibility to externally protect user data.

To protect user privacy within the myVisitPlanner processes, we first identified the necessary information that needs to be recorded in the user profiles, following the principle of data minimization. Then, we divided the data items into categories and identified the processes within myVisitPlanner that needed access to each data category. Finally we proposed methods to confine access to data categories only to the necessary processes.

8.2 Profiles in myVisitPlanner

In myVisitPlanner a usage profile in maintained and dynamically updated for each user to achieve effective service personalisation. This usage profile contains:

- User demographic characteristics: Age group, sex and disability status.
- Activity type preferences: Tuples in the form (Activity Type, Preference Value, Duration Preference, Time of Day Preference)
- System preferences: Activities duration preferences, free time preferences, etc.
- History information: Full selected travel plans, Activity ratings within selected plans

Activities in myVisitPlanner are categorised according to an activity type ontology. During profile creation users input their preferences regarding the available activity types. Each preference record is in the form (Activity Type, Activity Type Rating, Duration Preference, Time-of-day Preference). Additionally, in the history of a profile, the previously created travel plans, as well as the activity ratings given by the user are recorded.

To achieve usability and privacy-protection for the profile data, the profile history is sectioned into two parts. In the first part is the detailed usage history, where all interactions of the user to the system are recorded. The second part is the activity ratings, which result either directly from the user or indirectly from the user behaviour within the system. The activity ratings are used as input in the recommendation system algorithm and have the format (Profile, Activity, Rating).

8.3 Profile Data Use in myVisitPlanner

Profile data use in myVisitPlanner is divided into two categories:

1. The data use during the user interaction with the system. In this case, the profile data is used to provide the service to the user through the interface, for example, the profile editing functionality. In this case it is necessary for the data to be used in its original form.

2. The data use during the machine learning processes for the improvement of the recommendation system effectiveness. In this case the profile data can be used in anonymised form.

The processes in myVisitPlanner that handle profile information are the following:

- Profile editing, available during user interaction with the user interface.

- Item based recommendation. During this process, activities that have similar attributes as previously rated activities are recommended to the user. This process uses as input the user activity ratings in the form (Profile, Activity, Rating).

- User clustering, during which myVisitPlanner users are grouped according to their characteristics (demographic data and activity type preferences).

- Cluster-based recommendation, produces activity recommendations based on user clusters and the activity ratings from other users within the cluster.

- Scheduling. This process plans the selected recommended activities into a timetable, according to the user system preferences and the relevant activity characteristics (e.g. open hours).

In Table 8.1, the profile data categories are shown against the main processes where they are accessed and the entities that need access to the data.

8.4 Privacy Protection in myVisitPlanner

Having determined the profile data usage scope within the myVisitPlanner processes, applicable privacy-protection methods are proposed, that allow the effective operation of the processes.

Entity	User	Recommendation			Scheduler
Scope / Data Item	Profile Editing (UI)	Activity Similarity Based Recommendation	User Clustering	User Cluster Based Recommendation	Scheduling
Demographic Data	■		■		
Activity Type Preferences (in User Profile)	■		■		
System Preferences (in User Profile)	■				■
Detailed User Interaction Log	■				
Activity Ratings	■	■	■		

Table 8.1: Data use in myVisitPlanner processes

8.4.1 Encrypted User Profiles

As illustrated in Table 8.1, access to the raw profile user data is not needed when the user is offline. Therefore, user profiles are stored encrypted with the user password, prohibiting unauthorised access to them while the user is offline. Profile data is transparently decrypted whenever the user logs into the system, and is kept decrypted for the duration of the user's session and then re-encrypted automatically. This way, the time user data remain unencrypted, as well as the amount of unencrypted profiles on the system at each moment is minimized, along with the extent of possible data leak, in case of a security breach.

For the implementation of the profile data encryption and decryption functionality the OWASP Enterprise Security API (ESAPI) [1] Java open source library was used. The data in the database is encrypted using the AES128 symmetric encryption algorithm. The symmetric key is itself encrypted using another cipher, using the KEK (Key Encryption Key) scheme [95], to allow changing user encryption keys without needing to decrypt the data and re-encrypt with the new key.

Profile data extraction for use in the recommendation system processes

The recommendation system in myVisitPlanner contains processes that are executed when the user is offline, to improve the quality of the recommendations and need access to part of the user profile data. For this purpose, the needed information is extracted to separate database tables, when the user is online. Specifically, before the user logs off and the profile data is encrypted for storage, the activity ratings contained in the profile are extracted in the form (Profile,

[1] https://www.owasp.org/index.php/Category:OWASP_Enterprise_Security_API

Activity, Rating), as well as the activity type ratings are extracted in the form (Profile, ActivityType, Rating) and stored into separate tables in the database. The storage of the Profile field in the unencrypted table records is needed for the recommendation system processes, however it does not constitute a direct privacy threat, since only the ID is used while the rest of the profile is encrypted and its data cannot be used to identify the user.

8.4.2 Privacy Protection during User Clustering

The user clustering process divides users into groups according to their activity type preferences and their activity ratings and their demographic information. In order to protect user privacy, it is important to retain the k-anonymity property for the clusters, i.e. that all clusters contain a minimum amount of users (k), so that no individual users can be singled out from their clusters. For this purpose, an appropriate 3-stage clustering methodology was employed, that ensures compliance with the k-anonymity principle. This process may result into a less fine-grained user clustering, dividing users into a smaller number of groups, however it is considered a reasonable trade-off, since this measure reduces the effectiveness of possible de-anonymisation attacks.

8.4.3 Privacy Protection during Cluster-based Recommendations

To produce cluster-based activity recommendations for a user, the relevant activities that other users within the same cluster have previously rated are used. To protect user privacy during this process, data aggregation is used. Activity ratings of the cluster members are aggregated into one representative cluster user (centroid) to be used for generating recommendations. Moreover, system default profiles are created for certain user groups, in order to boost recommendations for certain activities and at the same time distort cluster data to further inhibit the effectiveness of possible attacks.

8.4.4 System Use Without Registration

Apart from the privacy protection techniques proposed in the previous sections, special care was taken to enable system use without registration, with as minimum as possible reduction of service capabilities. In this scenario users can create temporary profiles, in order to navigate through the system and create travel plans. Temporary profiles only contain demographic and user preference information, therefore it is not possible to implement recommendation and machine learning techniques based on the profile history. Nevertheless, the tempo-

rary profile can be used to categorise the user into a cluster, in order to receive cluster-based recommendations.

The ability to use myVisitPlanner without registration constitutes an important application feature, which demonstrates the privacy-respecting principles of its creation and the non-discriminating attitude towards users who choose not to provide any of their information for storage.

8.5 Discussion

The myVisitPlanner system is a real-world application, that aims to offer activity recommendations that help visitors and residents of the region of Northern Greece to plan their activities during their stay in this area. In this application user data is stored centrally, at the application servers, however privacy-protection mechanisms are proposed to minimize potential data loss from unauthorised access to user information. The proposed solutions do not achieve complete user data protection, however, they constitute a realistic and practical approach that was put into practice within the existing available infrastructure of participating companies and institutions.

Chapter 9

Conclusions, Challenges and Directions

Personal privacy during physical and electronic transactions is not being adequately protected today. Profiles containing individuals' personal information are created and stored outside the information owners' control. All facets of online activity today result in user personal data being recorded: online purchases, web searches, location services, social networking and many more applications lead to personal data disclosure and potential leaks. Therefore, it is important that applications today are created with privacy and security as built-in, mandatory features. The privacy-enhancing mechanisms needed to implement such applications exist today and can lead to a privacy-friendly digital world.

In order for this to happen, apart from the theoretical research needed to improve the effectiveness and efficiency of these tools, applied research is needed to put the available state-of-the-art technologies into practice. New solutions and applications need to be developed to create better privacy-preserving conditions for individuals and the industry. These solutions need to be designed to be integrated into current practices as transparently as possible, ensuring usability for the individuals and easy deployability for the institutions. To achieve this level of usability and ensure adoption by the industry a lot more work needs to be done, especially on the practical level, to provide effective solutions for real-world applications. Privacy-protection needs to be introduced into all kinds of electronic activities, addressing application-specific requirements and obstacles and demonstrating the feasibility of privacy-preserving approaches. This work makes a few steps in this direction.

In this work we created several privacy-preserving applications, following a common principle: Data owners should keep control of their information and utilize it to their own benefit under well-defined conditions. We showed that, utilizing contemporary cryptographic mechanisms, it is feasible to build practical

Privacy-preserving applications	Protected data items			Functionality / privacy aspects				
	PII and user attributes	Transaction History	Television Viewership Records	User Initiated Transactions	Controlled Access to data consumers	Privacy Economics	Financial compensation	Data Accountability/ Validation
Polis	■			■				
FPIT	■				■	■	■	
PrivTAM	■		■			■	■	■
Portfolio	■	■			■	■		■
PrivCAPTCHA	■			■				■

Table 9.1: Summary of privacy-preserving applications and their characteristics

applications that follow this principle. The Polis personal data management system proposed the main user-centric functionality during online purchases. The FPIT market for personal information, introduced the idea of providing controlled access to user information to interested parties, outside the context of online purchases, in exchange for fair compensation. The Portfolio consumer profile management system took the market idea one step further, proposing user-centric utilization of consumer profiles. The main ideas of FPIT and Portfolio were also implemented into practical applications. The FPIT idea was introduced into PrivTAM, a privacy-preserving Television Audience Measurement system, while the Portfolio idea was applied to create PrivCATCHA, a privacy-preserving VoIP CAPTCHA adaptation system. The main characteristics of the privacy-preserving applications developed within the context of this dissertation are summarised in Table 9.1. Additionally, a more applied methodology for user profile privacy protection was proposed and implemented into myVisitPlanner, a real-world travel recommendation and planning system.

We expect the proposed solutions to appeal both to individuals and companies. Adopting the proposed applications, individuals are able to interact with retailers and service providers and receive personalised services without having to give up control of their personal information. Additionally, security and privacy protection issues are simplified for companies, which are able to provide their services without having to maintain and protect customer profiles at their servers. As privacy awareness grows, individuals are more likely to prefer doing business with companies that adopt privacy-preserving solutions. Therefore privacy protection can act as a market advantage for companies. Moreover, data processing entities are able to legally gain access to accountable, up-to-date consumer information, leading to better quality results and fewer privacy risks.

Possible future directions of our work include the examination of new application scenarios, within the premise of user-side data management, and the design of corresponding privacy-preserving solutions. Some possible interesting applications would be privacy-preserving public opinion polls, as well as health data management. A very interesting project would also be the integration of the proposed application-specific solutions into a unified application that would be able

to handle personal data protection and management regardless of the particular performed task. An additional possibility for future work is the development of practical solutions for personal data management in mobile devices.

Apart from the technical work, there is a need for across-the-board awareness of the availability and the potential of state-of-the-art privacy enhancing technologies and a demand for their use in contemporary applications. Application designers and developers need to be familiar with the recent technological advancements and the ways these can be employed. Furthermore, policy makers also need to be aware of the potential of the new technologies in order to introduce them into updated laws and regulations concerning privacy protection. Finally, public awareness of the privacy risks of current technologies and applications is important, as well as a demand for privacy-preserving applications, which will in turn motivate the industry to adopt privacy-protecting practices. Moreover, individuals need to realise the economic value of their information and adopt applications that acknowledge this value and compensate them for the use of their data.

Another important issue regarding privacy protection, is the existence of up-to-date laws and regulations that not only provide guidance and directions for privacy protection, but also have enforcement power, to safeguard privacy-protection and ensure concrete consequences for non-complying entities.

Given that privacy-preserving mechanisms require a more active handling of personal data at the user side, personal data management tools need to be developed with usability in mind. Ideally, a generalised user-friendly application would assist individuals in the management of their data, allowing them to control the disclosure and uses of their information in all relevant activities. For such an application to be developed, new standards regarding support for privacy-preserving operations and communications need to be created, along with corresponding programming interfaces (APIs) for the privacy-enhancing technologies. This way privacy-supporting applications can be uniformly developed and deployed.

In conclusion, we are yet far from adequately achieving overall privacy protection in the digital world, however important technological advancements in the past decades have allowed for important steps to be made toward this direction. Privacy-protection can be embedded into the design of contemporary applications and what remains is for specific solutions to be developed to overcome the practical obstacles and put privacy enhancing technologies into practice.

References

[1] M. Ackerman. The intellectual challenge of cscw: The gap between social requirements and technical feasibility, 2000. 32

[2] A. Acquisti. Privacy and security of personal information: Technological solutions and economic incentives. In J. Camp and R. Lewis, editors, *The Economics of Information Security*, pages 165–178. Kluwer, 2004. 2, 18, 22, 51

[3] A. Acquisti, S. Gritzalis, C. Lambrinoudakis, and S. De Capitani di Vimercati. *Digital privacy.* Auerbach Publications, Taylor & Francis Group, 2008. 61

[4] A. Acquisti. Privacy and security of personal information. In *Economics of Information Security*, pages 179–186. Springer, 2004. 10

[5] E. Adar and B. A. Huberman. A market for secrets. *First Monday*, **6**:200–1, 2001. 36

[6] G. Aggarwal, M. Bawa, P. Ganesan, H. Garcia-Molina, K. Kenthapadi, R. Motwani, U. Srivastava, D. Thomas, and Y. X. 0002. Two can keep a secret: A distributed architecture for secure database services. In *CIDR*, pages 186–199, 2005. 25

[7] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu. Hippocratic databases. In *VLDB '2002: Proceedings of the 28th international conference on Very Large Data Bases*, pages 143–154. VLDB Endowment, 2002. 21

[8] F. Alvarez, C. Martin, D. Alliez, P. Roc, P. Steckel, J. Menendez, G. Cisneros, and S. Jones. Audience measurement modeling for convergent broadcasting and iptv networks. *Broadcasting, IEEE Transactions on*, **55**[2]:502–515, June 2009. 50, 51

[9] R. Anderson. U.K. government loses personal data on 25 million citizens. *EDRI-gram*, **Number 5.22**, 21 November 2007. 1, 23

[10] R. Anderson, C. Manifavas, and C. Sutherland. Netcard – a practical electronic-cash system. pages 49–57. 1997. 40

[11] R. Anderson and T. Moore. The economics of information security. *Science*, **314**[5799]:610–613, 2006. 10

[12] C. Ardagna, J. Camenisch, M. Kohlweiss, R. Leenes, G. Neven, B. Priem, P. Samarati, D. Sommer, and M. Verdicchio. Exploiting cryptography for privacy-enhanced access control: A result of the PRIME project. *Journal of Computer Security*, **18**[1]:123–160, 2010. 15

[13] G. Ateniese, D. Chou, B. de Medeiros, and G. Tsudik. Sanitizable signatures. In S. di Vimercati, P. Syverson, and D. Gollmann, editors, *Computer Security—ESORICS 2005*, **3679** of *Lecture Notes in Computer Science*, pages 159–177. Springer Berlin Heidelberg, 2005. 14, 85

[14] V. Ayala-Rivera, P. McDonagh, T. Cerqueus, and L. Murphy. A systematic comparison and evaluation of k-anonymization algorithms for practitioners. *Transactions on Data Privacy*, **7**[3]:337–370, 2014. 12

[15] F. Baldimtsi and A. Lysyanskaya. Anonymous credentials light. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, CCS '13, pages 1087–1098, New York, NY, USA, 2013. ACM. 70

[16] T. Balopoulos, S. Gritzalis, and S. Katsikas. Specifying and implementing privacy-preserving cryptographic protocols. *International Journal of Information Security*, **7**[6]:395–420, 2008. 63

[17] E. Bangerter, J. Camenisch, and A. Lysyanskaya. A cryptographic framework for the controlled release of certified data. In B. Christianson, B. Crispo, J. Malcolm, and M. Roe, editors, *Security Protocols Workshop*, **3957** of *LNCS*, pages 20–42. Springer, 2004. 18

[18] O. Baudron, P.-A. Fouque, D. Pointcheval, J. Stern, and G. Poupard. Practical multi-candidate election system. In *Proceedings of the twentieth annual ACM symposium on Principles of distributed computing*, PODC '01, pages 274–283, New York, NY, USA, 2001. ACM. 49, 55, 57, 62, 69, 70, 71

[19] P. Bichsel and J. Camenisch. Mixing identities with ease. In E. Leeuw, S. Fischer-Hübner, and L. Fritsch, editors, *Policies and Research in Identity Management*, **343** of *IFIP Advances in Information*

and Communication Technology, pages 1–17. Springer Berlin Heidelberg, 2010. 103

[20] J. P. Bigham and A. C. Cavender. Evaluating existing audio captchas and an interface optimized for non-visual use. In *Proceedings of the 27th international conference on Human factors in computing systems*, pages 1829–1838, Boston, MA, USA, 2009. ACM. 94

[21] P. Bogetoft, D. L. Christensen, I. Damgård, M. Geisler, T. Jakobsen, M. Krøigaard, J. D. Nielsen, J. B. Nielsen, K. Nielsen, J. Pagter, M. Schwartzbach, and T. Toft. Secure multiparty computation goes live. In R. Dingledine and P. Golle, editors, *Financial Cryptography and Data Security*, pages 325–343. Springer-Verlag, Berlin, Heidelberg, 2009. 49

[22] K. Bohrer and B. Holland, editors. *Customer Profile Exchange (CPExchange) Specification*. IDEAlliance, 2000. http://www.idealliance.org/cpexchange. 20

[23] C. Bösch, P. Hartel, W. Jonker, and A. Peter. A survey of provably secure searchable encryption. *ACM Comput. Surv.*, **47**[2]:18:1–18:51, August 2014. 13

[24] E. Bursztein, S. Bethard, C. Fabry, J. Mitchell, and D. Jurafsky. How good are humans at solving captchas? a large scale evaluation. In *Proceedings of the 2010 IEEE Symposium on Security and Privacy*, pages 399–413, Oakland, California, USA, 2010. IEEE. 94

[25] J. Camenisch. Information privacy?! *Computer networks*, **56**[18]:3834–3848, 2012. 9, 10, 11

[26] J. Camenisch, M. Dubovitskaya, M. Kohlweiss, J. Lapon, and G. Neven. Cryptographic mechanisms for privacy. In J. Camenisch, S. Fischer-Hübner, and K. Rannenberg, editors, *Privacy and Identity Management for Life*, pages 117–134. Springer, Berlin / Heidelberg, 2011. 98, 102

[27] J. Camenisch et al. Specification of the identity mixer cryptographic library, version 2.3.4, february 10, 2012. 51

[28] J. Camenisch and T. Gross. Efficient attributes for anonymous credentials, 2008. 103

[29] J. Camenisch, T. Gross, and T. Heydt-Benjamin. Accountable privacy supporting services. *Identity in the Information Society*, **2**[3]:241–267, December 2009. 2

[30] J. Camenisch, M. Koprowski, and B. Warinschi. Efficient blind signatures without random oracles. In C. Blundo and S. Cimato, editors, *Security in Communication Networks*, **3352** of *Lecture Notes in Computer Science*, pages 134–148. Springer Berlin Heidelberg, 2005. 51, 64, 65, 70

[31] J. Camenisch and A. Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In B. Pfitzmann, editor, *Advances in Cryptology — EUROCRYPT 2001*, **2045** of *Lecture Notes in Computer Science*, pages 93–118. Springer Berlin Heidelberg, 2001. 12, 51, 64, 70, 103

[32] J. Camenisch and A. Lysyanskaya. A signature scheme with efficient protocols. In S. Cimato, G. Persiano, and C. Galdi, editors, *Security in Communication Networks*, **2576** of *Lecture Notes in Computer Science*, pages 268–289. Springer Berlin / Heidelberg, 2003. 12

[33] J. Camenisch and B. Pfitzmann. Federated identity management. In M. Petković and W. Jonker, editors, *Security, Privacy, and Trust in Modern Data Management*, Data-Centric Systems and Applications, pages 213–238. Springer, Berlin / Heidelberg, 2007. 12, 96, 98, 100

[34] D. Cash, J. Jaeger, S. Jarecki, C. Jutla, H. Krawczyk, M.-C. Rosu, and M. Steiner. Dynamic searchable encryption in very-large databases: Data structures and implementation. In *Network and Distributed System Security Symposium (NDSS'14)*, 2014. 14, 84, 85

[35] D. Cash, S. Jarecki, C. Jutla, H. Krawczyk, M.-C. Roşu, and M. Steiner. Highly-scalable searchable symmetric encryption with support for boolean queries. In R. Canetti and J. Garay, editors, *Advances in Cryptology – CRYPTO 2013*, **8042** of *Lecture Notes in Computer Science*, pages 353–373. Springer Berlin Heidelberg, 2013. 14, 84, 85, 91

[36] S.-C. Cha and Y.-J. Joung. From p3p to data licenses. In *Privacy Enhancing Technologies*, pages 205–222, 2003. 11, 21, 25

[37] D. Chaum. Blind signatures for untraceable payments. In D. Chaum, R. Rivest, and A. Sherman, editors, *Advances in Cryptology*, pages 199–203. Springer US, 1983. 51, 64, 65, 70

[38] V. Ciriani, S. Capitani di Vimercati, S. Foresti, and P. Samarati. κ-anonymity. In *Secure Data Management in Decentralized Systems*, **33** of *Advances in Information Security*, pages 323–353. Springer, 2007. 62

[39] ConsumerReports. C.R. investigates: Your privacy for sale. *Consumer Reports*, **71**[10]:41, October 2006. http://www.accessmylibrary.com/coms2/summary_0286-29062087_ITM. 23

[40] J. Crosby. Challenges and opportunities in identity assurance. Technical report, HM Treasury, United Kingdom, March 2008. http://www.hm-treasury.gov.uk/d/identity_assurance060308.pdf. 3

[41] X. Dai, J. Grundy, and B. W. N. Lo. Comparing and contrasting micro-payment models for ecommerce systems. In *International Conferences of Info-tech and Info-net (ICII)*, **47**, 2001. 40

[42] I. Damgård, M. Jurik, and J. Nielsen. A generalization of paillier's public-key system with applications to electronic voting. *International Journal of Information Security*, **9**[6]:371–385, 2010. 49

[43] I. Damgård and M. Jurik. A generalisation, a simplification and some applications of paillier's probabilistic public-key system. In *Proceedings of the 4th International Workshop on Practice and Theory in Public Key Cryptography: Public Key Cryptography*, PKC '01, pages 119–136, London, UK, 2001. Springer-Verlag. 55, 72

[44] G. Danezis, J. Domingo-Ferrer, M. Hansen, J.-H. Hoepman, D. L. Metayer, R. Tirtea, and S. Schiffner. Privacy and data protection by design-from policy to engineering. *arXiv preprint arXiv:1501.03726*, 2015. 9, 12

[45] R. Dantu, S. Fahmy, H. Schulzrinne, and J. Cangussu. Issues and challenges in securing voip. *Computers & Security*, **28**[8]:743–753, 2009. 109

[46] Y.-A. de Montjoye, E. Shmueli, S. S. Wang, and A. S. Pentland. openpds: Protecting the privacy of metadata through safeanswers. *PLoS ONE*, **9**[7], 2014. 15

[47] Y.-A. de Montjoye, S. S. Wang, A. Pentland, D. T. T. Anh, A. Datta, et al. On the trusted use of large-scale personal data. *IEEE Data Eng. Bull.*, **35**[4]:5–8, 2012. 15

[48] Y. Deswarte and S. Gambs. A proposal for a privacy-preserving national identity card. *Transactions on Data Privacy*, **3**[3]:253–276, 2010. 15, 96

[49] C. Diaz and B. Preneel. Accountable anonymous communication. In M. Petković and W. Jonker, editors, *Security, Privacy, and Trust in Modern Data Management*, Data-Centric Systems and Applications, pages 239–253. Springer, Berlin / Heidelberg, 2007. 96

[50] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, pages 303–320, August 2004. 27, 28

[51] DISCREET. Discreet service provision in smart environments. FP6-2004-IST-4 contract no. 27679, 2008. http://www.ist-discreet.org/. 20

[52] G. Drosatos, P. S. Efraimidis, A. Arampatzis, G. Stamatelatos, and I. N. Athanasiadis. Pythia: A privacy-enhanced personalized contextual suggestion system for tourism. In *39th Annual IEEE Computers, Software and Applications Conference (COMPSAC 2015)*, pages 822–827. IEEE Computer Society, july 2015. 15

[53] G. Drosatos and P. Efraimidis. Privacy-preserving statistical analysis on ubiquitous health data. In S. Furnell, C. Lambrinoudakis, and G. Pernul, editors, *Trust, Privacy and Security in Digital Business (TrustBus '11)*, **6863** of *LNCS*, pages 24–36. Springer Berlin Heidelberg, 2011. 53

[54] D. Eastlake 3rd. Publicly Verifiable Nominations Committee (NomCom) Random Selection. RFC 3797 (Informational), June 2004. 53, 55

[55] P. S. Efraimidis, G. Drosatos, F. Nalbadis, and A. Tasidou. Towards privacy in personal data management. *Information Management and Computer Security (IMCS)*, **17**[4]:311 – 329, 2009. 17

[56] P. S. Efraimidis, G. Drosatos, F. Nalbadis, and A. Tasidou. Towards privacy in personal data management. In *Proceedings of the 12th Pan-Hellenic Conference on Informatics (PCI '08)*, pages 3–7, August 2008. 17, 37, 43

[57] S. Ehlert, D. Geneiatakis, and T. Magedanz. Survey of network security systems to counter sip-based denial-of-service attacks. *Computers & Security*, **29**[2]:225–243, 2010. 110

[58] European Parliament. Directive 95/46/EC. In *Official Journal L 281*, pages 0031–0050. 24 October 1995. http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31995L0046:EN:HTML. 9

[59] M. Fahrmair, W. Sitou, and B. Spanfelner. Security and privacy rights management for mobile and ubiquitous computing. In *Workshop on UbiComp Privacy*, 2005. 21

[60] S. Fischer-Hübner. *IT-security and Privacy: Design and Use of Privacy-enhancing Security Mechanisms*. Springer-Verlag, Berlin, Heidelberg, 2001. 9

[61] S. Gambs, S. Ranellucci, and A. Tapp. The crypto-democracy and the trustworthy (position paper). In J. Garcia-Alfaro, J. Herrera-JoancomartΓ, E. Lupu, J. Posegga, A. Aldini, F. Martinelli, and N. Suri, editors, *Data Privacy Management, Autonomous Spontaneous Security, and Security Assurance*, **8872** of *Lecture Notes in Computer Science*, pages 58–76. Springer International Publishing, 2015. 51

[62] D. Geneiatakis, G. Kambourakis, C. Lambrinoudakis, T. Dagiuklas, and S. Gritzalis. A framework for protecting a SIP-based infrastructure against malformed message attacks. *Computer Networks*, **51**[10]:2580–2593, 2007. 110

[63] C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the 41st annual ACM symposium on Theory of computing (STOC '09)*, pages 169–178, New York, NY, USA, 2009. ACM. 59

[64] G. Georgakopoulos. *Privacy enhancing technologies for personal data management*. Master's thesis, Dept. Electrical and Computer Engineering, Democritus Univ. of Thrace, Greece, October 2008. 43

[65] I. Goldberg. *A Pseudonymous Communications Infrastructure for the Internet*. PhD thesis, Univ. of California at Berkeley, 2000. 18

[66] I. Goldberg. Privacy-enhancing technologies for the internet III: Ten years later. In A. Acquisti, S. Gritzalis, C. Lambrinoudakis, and S. di Vimercati, editors, *Chapter 1 of Digital Privacy: Theory, Technologies, and Practices*. December 2007. 1, 18, 27

[67] O. Goldreich. *The Foundations of Cryptography*, **2**. Cambridge University Press, 2004. 61

[68] O. Goldreich and Y. Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, **7**[1]:1–32, 1994. 51, 64

[69] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof-systems. In *Proceedings of the seventeenth Annual ACM Symposium on Theory of computing*, STOC '85, pages 291–304, New York, NY, USA, 1985. ACM. 51, 64

[70] S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, **28**[2]:270 – 299, 1984. 61

[71] P. Golle, F. McSherry, and I. Mironov. Data collection with self-enforcing privacy. In *CCS '06: 13th ACM conference on Computer and communications security*, pages 69–78, New York, NY, USA, 2006. ACM. 25

[72] R. Gopal, R. Garfinkel, M. Nunez, and D. Rice. Electronic markets for private information: Economic and security considerations. In *HICSS '06: Proceedings of the 39th Annual Hawaii International Conference on System Sciences*, Washington, DC, USA, 2006. IEEE Computer Society. 36

[73] R. Greenstadt and M. D. Smith. Protecting personal information: Obstacles and directions. In *Fourth Workshop on the Economics of Information Security (WEIS05) 2005*, 2005. 10

[74] D. Gritzalis. Embedding privacy in IT applications development. *Inf. Manag. Comput. Security*, **12**[1]:8–26, 2004. 18

[75] D. Gritzalis, K. Moulinos, and K. Kostis. A privacy-enhancing e-business model based on infomediaries. In *Proceedings of the International Workshop on Information Assurance in Computer Networks: Methods, Models, and Architectures for Network Security*, MMM-ACNS '01, pages 72–83, London, UK, UK, 2001. Springer-Verlag. 18

[76] S. Gritzalis. Enhancing web privacy and anonymity in the digital era. *Information Management and Computer Security*, **12**[3]:255–287, 2004. 10, 18

[77] J. Hong. *An Architecture for Privacy-Sensitive Ubiquitous Computing*. PhD thesis, University of California at Berkeley, Computer Science Division, Berkeley, 2005. 18, 25

[78] J. Jang-Jaccard and S. Nepal. A survey of emerging threats in cybersecurity. *Journal of Computer and System Sciences*, **80**[5]:973 – 993, 2014. 48

[79] P. Jäppinen. *ME - Mobile Electronic Personality*. PhD thesis, Lappeenranta University of Technology, Finland, 2004. 18

[80] S. Jarecki, C. Jutla, H. Krawczyk, M. Rosu, and M. Steiner. Outsourced symmetric private information retrieval. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 875–888. ACM, 2013. 13, 14, 84, 85

[81] N. Jentzsch. *Theory of Information and Privacy*, pages 7–59. Springer, 2007. 32

[82] R. Johnson, D. Molnar, D. Song, and D. Wagner. Homomorphic signature schemes. In B. Preneel, editor, *Topics in Cryptology—CT-RSA 2002*, **2271** of *Lecture Notes in Computer Science*, pages 244–262. Springer Berlin Heidelberg, 2002. 14, 85

[83] A. Juels, M. Luby, and R. Ostrovsky. Security of blind digital signatures (extended abstract). In *Proceedings of the 17th Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '97, pages 150–164. Springer-Verlag, 1997. 70

[84] S. Kamara. Encrypted search. *XRDS: Crossroads*, **21**[3]:30–34, March 2015. 13

[85] G. Karjoth and M. Schunter. A privacy policy model for enterprises. In *Proceedings of the 15th IEEE workshop on Computer Security Foundations*, CSFW '02, pages 271–282, Washington, DC, USA, 2002. IEEE Computer Society. 20

[86] G. Karjoth, M. Schunter, and M. Waidner. Platform for enterprise privacy practices: privacy-enabled management of customer data. In *Proceedings of the 2nd international conference on Privacy enhancing technologies*, PET'02, pages 69–84, Berlin, Heidelberg, 2003. Springer-Verlag. 18, 20

[87] C. Karlof, N. Sastry, and D. Wagner. Cryptographic voting protocols: a systems perspective. In *Proceedings of the 14th conference on USENIX Security Symposium - Volume 14*, pages 33–50, Berkeley, CA, USA, 2005. USENIX Association. 49

[88] V. Katos and A. Patel. A partial equilibrium view on security and privacy. *Information Management & Computer Security*, **16**:74–83, 2008. 10, 51

[89] S. Katsikas, J. Lopez, and G. Pernul. Trust, privacy and security in e-business: Requirements and solutions. In *Panhellenic Conference on Informatics*, pages 548–558, 2005. 10, 26

[90] A. D. Keromytis. A comprehensive survey of voice over ip security research. *Communications Surveys & Tutorials, IEEE*, **14**[2]:514–537, 2012. 109

[91] L. Kissner and D. Song. Privacy-preserving set operations. In V. Shoup, editor, *Advances in Cryptology - CRYPTO '05*, **3621** of *Lecture Notes in Computer Science*, pages 241–257. Springer Berlin / Heidelberg, 2005. 61

[92] J. Kleinberg, C. Papadimitriou, and P. Raghavan. On the value of private information. In *Proceedings of the 8th conference on Theoretical aspects of rationality and knowledge*, pages 249–257. Morgan Kaufmann Publishers Inc., 2001. 10, 51

[93] L. Korba and S. Kenny. Towards meeting the privacy challenge: Adapting drm. In *Digital Rights Management (LNCS 2696/2003)*, pages 118–136. Springer Berlin / Heidelberg, 2003. 21

[94] S. Kremer, M. Ryan, and B. Smyth. Election verifiability in electronic voting protocols. In *Proceedings of the 15th European conference on Research in computer security*, ESORICS'10, pages 389–404, Berlin, Heidelberg, 2010. Springer-Verlag. 49

[95] P. Landrock. Key encryption key. In *Encyclopedia of Cryptography and Security*, pages 326–327. Springer, 2005. 114

[96] K. Laudon. Markets and privacy. *Commun. ACM*, **39**[9]:92–104, 1996. 17, 26, 36

[97] J. Lazar, J. Feng, T. Brooks, G. Melamed, B. Wentz, J. Holman, A. Olalere, and N. Ekedebe. The soundsright captcha: an improved approach to audio human interaction proofs for blind users, 2012. 95

[98] S. Lederer, J. Hong, A. Dey, and J. Landay. Personal privacy through understanding and action: Five pitfalls for designers. In *Designing Secure Systems That People Can Use*, pages 421–445. 2005. 22, 32

[99] H.-H. Lee and M. Stamp. An agent-based privacy-enhancing model. *Information Management & Computer Security*, **16**[3]:305–319, 2008. 18

[100] M. Lesk. Micropayments: An idea whose time has passed twice? *Security & Privacy, IEEE*, **2**[1]:61–63, 2004. 52

[101] Y. Lindell and B. Pinkas. Secure multiparty computation for privacy-preserving data mining. *Journal of Privacy and Confidentiality*, **1**:59–98, 2009. 54

[102] G. Lioudakis, E. Koutsoloukas, N. Dellas, N. Tselikas, S. Kapellaki, G. Prezerakos, D. Kaklamani, and I. Venieris. A middleware architecture for privacy protection. *Comput. Networks*, **51**[16]:4679–4696, 2007. 18

[103] M. M. Merener. Theoretical results on de-anonymization via linkage attacks. *Transactions on Data Privacy*, **5**[2]:377–402, 2012. 12

[104] S. Micali and R. L. Rivest. Micropayments revisited. In *Topics in Cryptology—CT-RSA 2002*, **2271**, pages 149–163. Springer, 2002. 40, 51

[105] L. Millett, B. Friedman, and E. Felten. Cookies and web browser design: toward realizing informed consent online. In *SIGCHI Conference on Human factors in computing systems*, pages 46–52, New York, USA, 2001. ACM. 22

[106] D. Mulligan and A. Schwartz. Your place or mine?: privacy concerns and solutions for server and client-side storage of personal information. In *Proceedings of the tenth conference on Computers, freedom and privacy: challenging the assumptions (CFP '00)*, pages 81–84, New York, NY, USA, 2000. ACM. 18, 24

[107] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Consulted*, **1**[2012]:28, 2008. 52

[108] T. Nishide and K. Sakurai. Distributed paillier cryptosystem without trusted dealer. In *Proceedings of the 11th international conference on Information security applications*, WISA'10, pages 44–60, Berlin, Heidelberg, 2011. Springer-Verlag. 52, 56, 61, 72

[109] A. Odlyzko. Privacy, economics, and price discrimination on the internet. In *Proceedings of the 5th international conference on Electronic commerce*, pages 355–366. ACM, 2003. 11

[110] A. Odlyzko. The case against micropayments. In *Financial Cryptography*, pages 77–83. Springer, 2003. 52

[111] P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *EUROCRYPT'99*, **1592** of *LNCS*, pages 223–238. Springer, 1999. 56, 58, 61

[112] L. Palen and P. Dourish. Unpacking "privacy" for a networked world. In *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 129–136, New York, NY, USA, 2003. ACM. 32

[113] J. Pieprzyk, T. Hardjono, and J. Seberry. *Fundamentals of computer security*, chapter 15. Monographs in theoretical computer science. Springer, 2003. 49

[114] A. Poller, U. Waldmann, S. Vowe, and S. Turpe. Electronic identity cards for user authentication; promise and practice. *IEEE Security & Privacy*, **10**[1]:46–54, 2012. 96

[115] T. Poutanen, H. Hinton, and M. Stumm. Netcents: a lightweight protocol for secure micropayments. In *WOEC'98: Proceedings of the 3rd conference on USENIX Workshop on Electronic Commerce*, page 3, Berkeley, CA, USA, 1998. USENIX Association. 40

[116] G. A. Pratt. Is a cambrian explosion coming for robotics? *The Journal of Economic Perspectives*, **29**[3]:51–60, 2015. 10

[117] PRIME. Privacy and identity management for europe. EC Contract No. IST-2002-507591, 2008. https://www.prime-project.eu/. 26

[118] PrimeLife. Bringing sustainable privacy and identity management to future networks and services. http://www.primelife.eu. 15

[119] E. Rahm and H. H. Do. Data cleaning: Problems and current approaches. *IEEE Bulletin of the Technical Committee on Data Engineering*, **23**[4], December 2000. 23

[120] I. Refanidis, C. Emmanouilidis, I. Sakellariou, A. Alexiadis, R.-A. Koutsiamanis, K. Agnantis, A. Tasidou, F. Kokkoras, and P. S. Efraimidis. myvisitplanner gr: Personalized itinerary planning system for tourism. In A. Likas, K. Blekas, and D. Kalles, editors, *Artificial Intelligence: Methods and Applications*, **8445** of *Lecture Notes in Computer Science*, pages 615–629. Springer International Publishing, 2014. 111

[121] Reuters. Axel springer hit by new german data leak scandal, 18th October 2008. http://www.reuters.com/article/internetNews/idUSTRE49H1GH20081018. 23

[122] R. Rivest, L. Adleman, and M. Dertouzos. On data banks and privacy homomorphisms. In *Foundations of Secure Computation*, pages 169–177. Academic Press, 1978. 59

[123] R. L. Rivest and A. Shamir. Payword and micromint: two simple micropayment schemes. In *CryptoBytes*, **2**, pages 69–87, 1996. 36, 40, 51, 60

[124] H. Rossnagel, J. Camenisch, L. Fritsch, T. Gross, D. Houdeau, D. Hühnlein, A. Lehmann, and J. Shamah. Futureid-shaping the future of electronic identity. *Datenschutz und Datensicherheit (DuD)*, **36**[3]:189–194, 2012. 15

[125] A. Sabouri, I. Krontiris, and K. Rannenberg. Attribute-based credentials for trust (abc4trust). In S. Fischer-Hübner, S. Katsikas, and G. Quirchmayr, editors, *Trust, Privacy and Security in Digital Business*, **7449** of *Lecture Notes in Computer Science*, pages 218–219. Springer Berlin Heidelberg, 2012. 15

[126] F. Salim, N. Sheppard, and R. Safavi-Naini. Enforcing p3p policies using a digital rights management system. In N. Borisov and P. Golle, editors, *Privacy Enhancing Technologies*, **4776** of *LNCS*, pages 200–217. Springer, 2007. 18

[127] P. Samuelson. Privacy as intellectual property? *Stanford Law Review*, **52**:1125, 2000. 17

[128] J. Shi, R. Zhang, Y. Liu, and Y. Zhang. PriSense: Privacy-preserving data aggregation in people-centric urban sensing systems. In *Proceedings of the 29th IEEE International Conference on Computer Communications (INFOCOM '10)*, pages 1–9. IEEE, March 2010. 51

[129] Y. Soupionis and D. Gritzalis. Aspf: Adaptive anti-spit policy-based framework. In *Availability, Reliability and Security (ARES), 2011 Sixth International Conference on*, pages 153–160, 2011. 110

[130] Y. Soupionis and D. Gritzalis. Audio captcha: Existing solutions assessment and a new implementation for voip telephony. *Computers & Security*, **29**[5]:603–618, 2010. 95

[131] A. Tasidou, P. S. Efraimidis, and V. Katos. Economics of personal data management: Fair personal information trades. In A. B. Sideridis and C. Z. Patrikakis, editors, *Next Generation Society. Technological*

and Legal Issues, **26**, chapter 14, pages 151–160. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010. 26

[132] A. Tasidou and P. S. Efraimidis. Using personal portfolios to manage customer data. In J. Garcia-Alfaro, G. Navarro-Arribas, N. Cuppens-Boulahia, and S. de Capitani di Vimercati, editors, *Data Privacy Management and Autonomous Spontaneus Security*, **7122** of *Lecture Notes in Computer Science*, pages 141–154. Springer Berlin Heidelberg, 2012. 78, 98

[133] A. Tasidou, P. S. Efraimidis, Y. Soupionis, L. Mitrou, and V. Katos. User-centric, privacy-preserving adaptation for voip captcha challenges. In *Proc. of the 6th International Symposium on Human Aspects of Information Security and Assurance (HAISA) 2012*, pages 139–149, 2012. 97

[134] U-Prove. http://www.credentica.com/. 15

[135] US Government. The cable tv privacy act of 1984. In *47 USC, Chapter 5, Subchapter V-A, Part IV, Sec. 551*. U.S. Gov. Printing Office, Washington, DC, 1984. 51

[136] US Government. Video privacy protection act. In *18 USC, Part I, Chapter 121, Sec. 2710, Pub.L. 100-618*. U.S. Gov. Printing Office, Washington, DC, 1988. 51

[137] UTD Data Security and Privacy Lab. Paillier threshold encryption toolbox, November 2011. 72

[138] R. A. Valdes-Benavides and P. L. Hernandez-Verme. Virtual currencies, micropayments and monetary policy: Where are we coming from and where does the industry stand? *Journal For Virtual Worlds Research*, **7**[3], 2014. 52

[139] H. Varian. Economic aspects of personal privacy. In *Privacy and self-regulation in the information age*. U.S. Dept. of Commerce, National Telecommunications and Information Administration, 1996. 10, 51

[140] P. Vassiliadis, A. Simitsis, and S. Skiadopoulos. Conceptual modeling for etl processes. In *5th ACM international workshop on Data Warehousing and OLAP*, pages 14–21, New York, USA, 2002. ACM. 23

[141] P. L. H. Verme and R. A. V. Benavides. Virtual currencies, micropayments and the payments systems: A challenge to fiat money and monetary policy? *European Scientific Journal*, **9**[19], 2013. 52

[142] W3C. The platform for privacy preferences 1.0 (p3p1.0) specification, 2002. http://www.w3.org/TR/P3P. 20

[143] D. J. Weitzner, H. Abelson, T. Berners-Lee, J. Feigenbaum, J. Hendler, and G. J. Sussman. Information accountability. *Commun. ACM*, **51**[6]:82–87, June 2008. 11

[144] Wikipedia. Fair trade. http://en.wikipedia.org/wiki/Fair_trade. 36

[145] D. Wilusz and J. Rykowski. The architecture of coupon-based, semi-off-line, anonymous micropayment system for internet of things. In L. Camarinha-Matos, S. Tomic, and P. Graça, editors, *Technological Innovation for the Internet of Things*, **394** of *IFIP Advances in Information and Communication Technology*, pages 125–132. Springer Berlin Heidelberg, 2013. 52

[146] J. Yan and A. S. E. Ahmad. Usability of captchas or usability issues in captcha design, 2008. 95

[147] Z. Yang, W. Lang, and Y. Tan. Fair micropayment system based on hash chains. *Tsinghua Science & Technology*, **10**[3]:328–333, June 2005. 40

[148] A. C. Yao. Protocols for secure computations. In *IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 160–164. IEEE, 1982. 49

[149] L.-C. Yeh, C.-S. Wang, C.-Y. Lin, and J.-S. Chen. An innovative application over communications-asa-service: Network-based multicast iptv audience measurement. In *Network Operations and Management Symposium (APNOMS), 2011 13th Asia-Pacific*, pages 1–7, Sept 2011. 50, 51

[150] J. H. Ziegeldorf, O. G. Morchon, and K. Wehrle. Privacy in the internet of things: threats and challenges. *Security and Communication Networks*, **7**[12]:2728–2742, 2014. 48