Computer Communications 36 (2013) 542–550

Contents lists available at SciVerse ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

Distributed multicast of fingerprinted content based on a rational peer-to-peer community

Josep Domingo-Ferrer [a], David Megías [b,*]

[a] Universitat Rovira i Virgili, UNESCO Chair in Data Privacy, Department of Computer Engineering and Mathematics, Av. Països Catalans 26, E-43007 Tarragona, Catalonia, Spain
[b] Universitat Oberta de Catalunya, Internet Interdisciplinary Institute (IN3), Estudis d'Informàtica, Multimèdia i Telecomunicació, Rambla del Poblenou, 156, E-08018 Barcelona, Catalonia, Spain

ARTICLE INFO

Article history:
Received 23 May 2012
Received in revised form 8 October 2012
Accepted 26 December 2012
Available online 4 January 2013

Keywords:
Co-utility
Multicast fingerprinting
Anonymous fingerprinting
Game theory

ABSTRACT

In conventional multicast transmission, one sender sends the same content to a set of receivers. This precludes fingerprinting the copy obtained by each receiver (in view of redistribution control and other applications). A straightforward alternative is for the sender to separately fingerprint and send in unicast one copy of the content for each receiver. This approach is not scalable and may implode the sender. We present a scalable solution for distributed multicast of fingerprinted content, in which receivers rationally co-operate in fingerprinting and spreading the content. Furthermore, fingerprinting can be anonymous, in order for honest receivers to stay anonymous.

1. Introduction

Copyright protection techniques have gained widespread attention by both academia and industry in the recent years. Home Internet access and the increased bandwidth of communications have contributed to the explosion of copyright-breaking copying of digital contents. In this context, fingerprinting emerged as a convenient technology to fight against unlawful digital content distribution [6,4].

Fingerprinting techniques consist of embedding a transparent watermark into the protected content in such a way that a unique identifier exists for each buyer of the content. This identifier can be extracted later on and might be used to trace and match an illegal distributor of the content. This makes it possible to undertake the appropriate legal actions against such treacherous buyers. Fingerprinting schemes can be classified in three different categories [7], namely symmetric, asymmetric and anonymous. In symmetric fingerprinting, the embedding of the fingerprint is performed by the merchant only and, thus, it provides no valid evidence of a treacherous behavior of a buyer (since the merchant herself could be the illegal distributor). In asymmetric fingerprinting, the embedding is performed using a protocol designed in such a way that only the buyer obtains the fingerprinted copy of the content. This makes it possible to prove the illegal distributor's treachery to a third party. Finally, anonymous fingerprinting retains the asymmetric property and also protects the privacy of buyers, whose identity is only revealed and disclosed in case of illegal distribution.

From the point of view of a buyer, anonymity is a valuable property and several protocols have been proposed for anonymous fingerprinting. However, current anonymous fingerprinting proposals in the literature (see Section 2.3 below for a brief review) place a substantial computational and communication burden on the merchant. The merchant's overhead is a relevant issue, since it will possibly result in buyer anonymity not being offered or offered at higher price by the merchant so that the latter can still enjoy some profit margin. Hence, the possibility of reducing the merchant's burden and the flexibility of choosing the watermarking technology freely among the best state-of-the-art techniques are worth investigating. This paper focuses on proposing a multicast approach to the anonymous fingerprinting problem which meets these two goals and shows a proof of concept with a practical implementation of the proposed system. The idea is to transfer the burden of a centralized fingerprinting technology to a distributed network of buyers who will collaborate to produce further copies of the fingerprinted contents.

Sending a content to N different receivers via multicast is much more bandwidth-efficient from the sender's point of view than performing N successive unicast transmissions. However, the unicast approach has the advantage of allowing the sender to fingerprint the content sent to each receiver. Unfortunately, the standard multicast approach does not allow fingerprinting: all

* Corresponding author.
E-mail addresses: josep.domingo@urv.cat (J. Domingo-Ferrer), dmegias@uoc.edu (D. Megías).

0140-3664/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.comcom.2012.12.005

receivers get exactly the same content. That is why a specific multicast anonymous fingerprinting protocol is proposed in this paper.

1.1. Contribution and plan of this paper

We specify a protocol whereby a sender manages to distribute a digital content to an unlimited number of receivers in such a way that:

- The content carries a different anonymous fingerprint for each receiver, so that unlawful content redistribution can be tracked; honest receivers stay anonymous.
- The sender does not need to fingerprint and send the content individually to each receiver; one fingerprinting and one unicast transmission by the server to one collaborative receiver are enough to bootstrap the process.
- Receivers are rationally interested to collaborate in forwarding and fingerprinting the content to other interested receivers (we call such rational collaboration co-utility); thanks to anonymous fingerprinting, intermediate receivers do not know the identities of the receivers they are forwarding the fingerprinted content to.

Section 2 gives some background on game theory, co-utility and anonymous fingerprinting. Section 3 describes the protocol and justifies its security. Section 4 argues the rational involvement by peers in game-theoretic terms and shows that our protocol achieves co-utility. Section 5 contains experimental results of a proof of concept. Section 6 summarizes conclusions and future research issues.

2. Background

2.1. Basics of game theory

A game is a protocol between a set of N *players*, $\{P^1, \ldots, P^N\}$. Each player P^i has her own *set of possible strategies*, say S_i. To play the game, each player P^i selects a strategy $s_i \in S_i$. We use $s = (s_1, \ldots, s_N)$ to denote the vector of strategies selected by the players and $S = \Pi_i S_i$ to denote the set of all possible ways in which players can pick strategies.

The vector of strategies $s \in S$ selected by the players determines the outcome for each player, which can be a payoff or a cost. In general, the outcome will be different for different players. To specify the game, we need to give, for each player, a preference ordering on these outcomes by giving a complete, transitive, reflexive binary relation on the set of all strategy vectors S. The simplest way to assign preferences is by assigning, for each player, a value for each outcome representing the payoff of the outcome (a negative payoff can be used to represent a cost). A function whereby player P^i assigns a payoff to each outcome is called a utility function and is denoted by $u_i : S \to \mathbb{R}$.

For a strategy vector $s \in S$, we use s_i to denote the strategy chosen by player P^i and s_{-i} to denote the $(N-1)$-dimensional vector of the strategies played by all other players. With this notation, the utility $u_i(s)$ can also be expressed as $u_i(s_i, s_{-i})$.

A strategy vector $s \in S$ is a *dominant strategy solution* if, for each player P^i and each alternate strategy vector $s' \in S$, it holds that

$$u_i(s_i, s'_{-i}) \geqslant u_i(s'_i, s'_{-i}) \tag{1}$$

In plain words, a dominant strategy s is the best strategy for each player P^i, independently of the strategies played by all other players.

A strategy vector $s \in S$ is said to be a *Nash equilibrium* if, for all players P^i and each alternate strategy $s'_i \in S_i$, it holds that

$$u_i(s_i, s_{-i}) \geqslant u_i(s'_i, s_{-i})$$

In plain words, no player P^i can change her chosen strategy from s_i to s'_i and thereby improve her payoff, assuming that all other players stick to the strategies they have chosen in s. A Nash equilibrium is self-enforcing in the sense that once the players are playing such a solution, it is in every player's best interest to stick to her strategy. Clearly, a dominant strategy solution is a Nash equilibrium. Moreover, if the solution is strictly dominant (*i.e.* when the inequality in Expression (1) is strict), it is also the unique Nash equilibrium. See [26] for further background on game theory.

2.2. Co-utility

We recall here the co-utility paradigm, which we introduced under the name general coprivacy in [16,17]. The following definition is simpler but equivalent to the one used in our previous papers.

Definition 1 (*Co-utility*). Let Π be a game with self-interested, rational players $P^1, \ldots, P^N$, with $N > 1$. Game Π is said to be *co-utile* with respect to the vector $U = (u_1, \ldots, u_N)$ of utility functions if there exist at least two players P^i and P^j, having strategies s^i and s^j, respectively, such that: (i) s^i involves P^i expecting co-operation from P^j; (ii) s^j involves P^j co-operating with P^i; (iii) (s^i, s^j) is an equilibrium for P^i and P^j in terms of u_i and u_j, respectively. In other words, there is co-utility between P^i and P^j, for some $1 \leqslant i, j \leqslant N$ with $i \neq j$, if the best strategy for P^i involves expecting co-operation from P^j and the best strategy for P^j is to co-operate.

If the equilibrium in Definition 1 is a Nash equilibrium, we have *Nash co-utility*. If the utility functions U in Definition 1 only consider privacy, co-utility becomes the plain coprivacy notion introduced in [16,17]; if utilities only consider security, we could speak of co-security; if they only consider functionality, co-utility becomes co-functionality.

2.3. Anonymous fingerprinting

Let $D_0 \in \{0,1\}^*$ denote some digital content (bit-string) some of whose bits can be changed in such a way that (i) the result remains "close" to D_0 (where "close" means "with a similar utility"), but (ii) without knowing which particular bits were changed, altering a "good portion" of these bits is impossible without rendering the content useless. The changed bits are usually called a mark or watermark; if bits are changed differently for each user receiving the content, the mark can also be called fingerprint. The algorithm used to embed a mark while satisfying the previous two conditions is called a watermarking algorithm; to embed a fingerprint can also be termed "to fingerprint". The second requirement above is actually the marking assumption stated in [6].

As mentioned in the introduction above, the type of fingerprinting relevant to our paper is anonymous fingerprinting. The first anonymous fingerprinting proposals relied on unspecified multiparty secure computation protocols [27,13]. In [14], an anonymous fingerprinting protocol completely specified from the computational point of view and based on committed oblivious transfers was presented. In [15], the tamper-proofness of a smart card on the buyer's side was used to simplify anonymous fingerprinting. More recent anonymous fingerprinting schemes rely on the homomorphic properties of public-key cryptography [20,30,21,22,25,29,28]. These schemes allow embedding the fingerprint in the encrypted domain. The buyer sends her encrypted fingerprint to the merchant who embeds it by operating with the encrypted content using the public key of the buyer. The resulting encrypted and fingerprinted content is sent to the buyer who can decrypt it using her private key. This way, only the buyer has access to the