SPIES: Secret Protection Incentive-based Escrow System

N. Boris Margolin Matthew K. Wright Brian N. Levine
Dept. of Computer Science, University of Massachusetts, Amherst, MA 01003
{margolin,mwright,brian}@cs.umass.edu

Abstract

Once electronic content has been released, it is very difficult to prevent perfect copies of the content from being widely distributed, which can cause economic harm to the content's owner and others. We focus on content which is to be shared to a limited extent, which is valuable, and which only needs to be protected for a limited amount of time, such as trade secrets. For such content we provide an economic incentive to limit sharing, without using DRM or watermarking. In our protocol, a quantity of money is placed in escrow, and anyone can get a portion of it by providing proof of knowledge of the content. Since payments become smaller as more individuals give proof, it is in the interest of those with access to the content to prevent further sharing.

1 Introduction

Content with economic value is being offered in digital form and transfered to collaborators, reviewers, distributors, and consumers over network connections. Digital content can easily be shared over peer-to-peer file sharing networks like Kazaa or Bittorrent. It is difficult for content owners to prevent dissemination of their works on those networks.

We present a protocol that provides an economic incentive to not share copies of digital content. Our scheme has the advantage of not being dependent on watermarking [9] or Digital Rights Management (DRM) software and hardware [7]. The protocol, called *Secret Protection Incentive-based Escrow System* (SPIES), is best suited for applications where the content must be protected for a limited period of time and shared among a limited set of persons.

Our intention is not to replace non-disclosure agreements, corporate policies and procedures, or other legal, technical, or physical protection layers. Rather, we aim to add economic incentives as an additional layer of protection.

SPIES has several stages. In the initialization phase, everyone who should have access to content places money into an escrow account. From then until the end of a distinct *protection period*, anyone who has a copy of the content can register anonymously with the escrow service. Registrants are entitled to receive a portion of the escrowed money once they prove they have the protected content. Therefore, legitimate possessors have an incentive to not share the file, lest they lose the money they have in escrow. It is even in the interest of unauthorized possessors not to propagate it further, as this will reduce the amount they could get from the escrow.

The remainder of this paper provides details of SPIES. We also discuss how different media can work with SPIES, including audio, video, and still images.

This work was supported in part by the National Science Foundation grants ANI-0087482, ANI-0133055, and EIA-0080119.

2 Related Work

A number of systems, such as Microsoft's Windows Media DRM [7], attempt to protect digital content from copying. Unlike such DRM systems, SPIES does not require special software or hardware to conceal content from the user. Indeed, the user can freely create backup copies in any format, even giving them to trusted third parties, since SPIES relies on the user's self-interest to control widespread sharing, rather than hiding the content.

Beyond DRM, Horne, Pinkas, and Sander [5] present the most similar work. In contrast to our approach, users are paid for sharing content with authorized users. The payments motivate users to keep content within a subscription community. It also differs from SPIES in that in that it requires a subscription community and is not self-limiting; users outside the subscription community who get access to the content have little incentive not to share the content widely. Golle, *et al* [4] study several incentive schemes to deal with the *free-rider* problem. Both of these works provide incentives for sharing information, rather than for keeping information private.

Brin, Davis, and Garcia-Molina [2] present a system for registering and checking for copied documents based on the same sentences appearing in both documents. Although the original idea was to combine this with enforcement, such a system could also be used with our incentives scheme, when the content being shared is a document.

SPIES makes use of anonymous communication. Fortunately, anonymity is a well-studied problem, and there are many existing protocols that provide sufficient protection. Working instantiations of these protocols include MixMinion [6], Tor [3], and the Anonymizer [1].

3 SPIES

The parties in SPIES (see Figure 1) are the legitimate possessors $A_1 \ldots A_k$ (where $k \geq 2$), an escrow service E, and a set of charities C that do not actively participate in the protocol. A_1 is the original possessor of the content. We denote the content as ϕ and time at which the content need no longer be protected as τ. The money which a legitimate possessor places in escrow with E is v. Finally, we denote any unauthorized possessors of the content (who may have obtained the content through theft or through unauthorized sharing) by $U_1 \ldots U_l$.

Variable	*Description*
ϕ	The content to be protected
$d(\phi)$	The description of ϕ
$H(\phi)$	The hash of ϕ
τ	End of period of content protection
A_1	Original possessor of ϕ
$A_1 \ldots A_k$	Authorized possessors of ϕ
$U_1 \ldots U_l$	Unauthorized possessors of ϕ
ρ	Total number $k + l$ of possessors
E	A trusted escrow service
C	A set of charities
c	One randomly-chosen charity from set C
$\$v$	Each A_i's monetary contribution
$\$kv$	The total amount of money in escrow
$f(\rho)$	The share size function, dependent on the number of possessors

Figure 1: Variables used in SPIES.

It is helpful to have a semantic description and a serial number for ϕ; we write this $d(\phi)$. We denote the exchange of x dollars as \$(x). The transfer of such funds can be done by any secure method, e.g., credit cards over SSL.[1]

Informally, all legitimate possessors A_i submit a payment $\$v$ to the escrow service E as a pledge that they will not share or sell the content before time τ. E keeps this money in trust until τ.

[1] For the sake of clarity, our details omit the fact that, where necessary, each party's message is signed for authenticity and integrity using previously setup up public or shared keys.

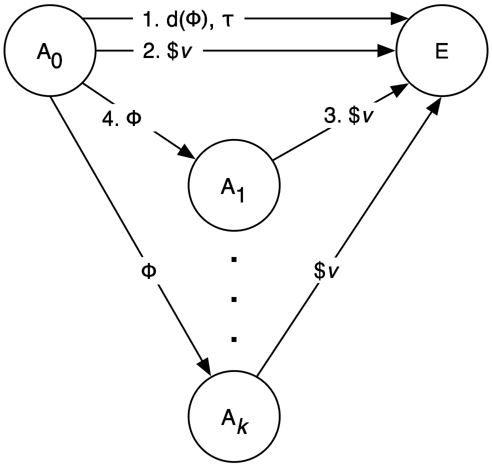

Figure 2: Phase 1: content is registered, monies are escrowed, and content is transfered.

Before time τ, anyone with knowledge of ϕ may anonymously submit a registration to E. After τ is reached, E will distribute a portion of the money in escrow to each registrant.

If A_i has given out copies of the content to an untrusted third party, U_j, then U_j may use that copy to register with E and reduce A_i's portion of money received from the escrow account. (It's important that no prosecution result from registering so as to encourage third parties to come forward.) So if A_i wants to maximize the return of money from escrow, she has an incentive to not share the content.

Phase 1: Exchange. A_1, the original possessor, registers a description of the content $d(\phi)$ and the ending time τ with the escrow agent E. All the participants A_i place $\$v$ in escrow, and A_1 sends ϕ to $A_2 \ldots A_k$. ϕ must be shared only with participants that send $\$v$ to E, but we do not require a specific mechanism for this; a protocol for shared secret exchange [8, pp. 122-124] is one possibility. At the end of this step, E holds $\$kv$ ($\$v$ from each A_i) and each of A_i has knowledge of the secret ϕ.

$$\begin{aligned} A_1 \rightarrow E &\;:\; d(\phi), \tau && (1)\\ A_1 \rightarrow E &\;:\; \$v && (2)\\ A_2 \rightarrow E &\;:\; \$v && (3)\\ A_1 \rightarrow A_2 &\;:\; \phi && (4)\\ &\;\vdots && \\ A_k \rightarrow E &\;:\; \$v && \\ A_1 \rightarrow E &\;:\; \phi && \end{aligned}$$

Phase 2: Registration. E publishes widely that it is seeking anonymous registrations from anyone holding content described by $d(\phi)$, including users U_l who have (illegitimately) obtained the content. Let ρ be the total number of registrants. As we will show, we expect to receive a single registration from each A_i and from each U_j, but the protocol remains effective no matter how many registrations are received.

First, A_1 sends a hash of the content to E, denoted as $H(\phi)$ below; being able to generate this hash will serve as proof of possession of ϕ. Each registrant anonymously[2] sends in the content description and a hash of the content. We discuss the details of H in Section 4. Although A_1 does not send in the hash again, she is considered one of the registrants. We denote the total number of registrations, $k + l$, by ρ; if no unauthorized sharing has occurred then $l = 0$ and $\rho = k$.

$$\begin{aligned} A_1 \rightarrow E &\;:\; H(\phi), d(\phi) && (5)\\ A_2 \rightarrow E &\;:\; H(\phi), d(\phi) && (6)\\ &\;\vdots && \\ A_k \rightarrow E &\;:\; H(\phi) d(\phi) && \\ U_1 \rightarrow E &\;:\; H(\phi), d(\phi) && (7)\\ &\;\vdots && \\ U_l \rightarrow E &\;:\; H(\phi), d(\phi) && \end{aligned}$$

Phase 3: Payment. At time τ, each registrant will get at most $1/\rho$ of the total amount held in escrow. If more than the expected number of registrations k occur, registrants will each get strictly *less* than $1/\rho$ of the total amount, specifically $1/(f(\rho)\rho)$.

$$E \rightarrow A_1 \;:\; \$\frac{kv}{f(\rho)\rho} \qquad (8)$$

[2] i.e. using a protocol such as Mixminion [6], Tor [3], or the Anonymizer [1].

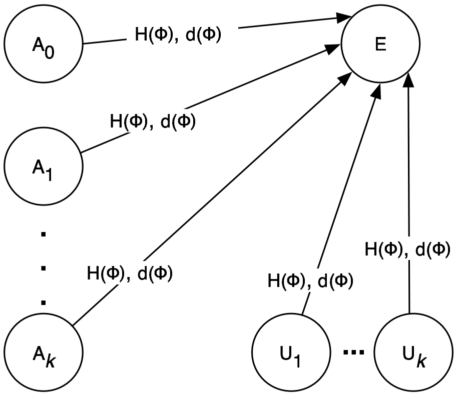

Figure 3: Phase 2: anonymous registration by possessors of the content.

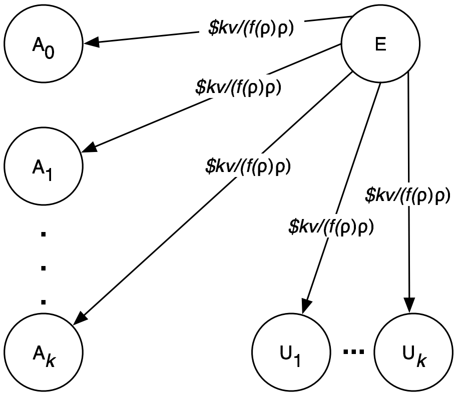

Figure 4: Phase 3: escrowed monies are distributed as $\$kv/(f(\rho)\rho)$ per person.

$$\begin{aligned} &\vdots \\ E \to A_k \quad &: \quad \$\frac{kv}{f(\rho)\rho} \\ E \to U_1 \quad &: \quad \$\frac{kv}{f(\rho)\rho} \\ &\vdots \\ E \to U_l \quad &: \quad \$\frac{kv}{f(\rho)\rho} \end{aligned} \tag{9}$$

Because of $f(\cdot)$, whenever more than the expected number (k) of registrations occur, the amount of money given to registrants will be less than the total in escrow. The remaining money is given to a randomly chosen charity $c \in C$, with the set C agreed upon between all A_i before the protocol. When a charity must be chosen, a secure random coin flip protocol (see Schneier [8] for examples) can be used among the originally authorized content holders. It is in the interest of content holders to collude with the charities; the large set of charities and secure coin flip is intended to make such collusion more difficult.

$$E \to c \quad : \quad \$kv\left(1 - \frac{1}{f(\rho)}\right) \tag{10}$$

3.1 The Payout Function f

We want a party adding additional dummy registrations to get strictly less total money. In the worst case, with only two registrants A and B, A may behave honestly and submit 1 registration while B submits k registrations, so that B receives $k/(k+1)$ of the shares. The size of the shares for the $k+1$ case must be decreased sufficiently such that the k shares together are strictly less than one of the larger shares.

To achieve this we construct a payout function $f(x) = 2^{x-k}$. Suppose that, before a participant B has registered, there are α existing registrations, and p dollars in escrow. B can choose any number of registrations β to add, to obtain β shares, each of size

$$\$\frac{p}{(\alpha+\beta)2^{\alpha+\beta-k}}$$

for a total payment of

$$\$\frac{\beta p}{(\alpha+\beta)2^{\alpha+\beta-k}}.$$

This is maximum at $\beta = 1$, i.e. when B registers exactly once.

When $\alpha + \beta = k$, which happens when there are only registrations from the authorized possessors, $f(\alpha + \beta) = 2^0 = 1$ and all the money in escrow is returned to the contributers. It is in B's interest to register exactly once, however many registrations k have already occurred.

Note that when someone is dishonest and submits multiple registrations, the other participants are harmed economically even more than the dishonest submitter; each share is exponentially smaller and they are only receiving one share each. This causes problems when there are large numbers of buyers, since the participants must trust more parties to be rational. Therefore, we do not suggest that the scheme will scale to the enormous number of buyers that, for example, would purchase a popular music album or movie.

3.2 Balancing Escrow with the Value of the Content

We can distinguish two types of problematic participants: those who allow others to access the content due to lax security measures, and those who want to make a monetary profit by selling the content without authorization. For the first type of participant, the escrow money serves an additional incentive to prevent access to the content and needs only to be an amount that is significant to the participants. For example, a reviewer of a well-known magazine may be trusted to not sell a review copy of a movie on the black market; the escrowed money gives her an incentive to be more careful with the content. Content producers already have trust mechanisms in place for releasing content, and SPIES can function as an additional layer of protection. In the second case, SPIES will only be effective when participants are willing to escrow more than what they could potentially earn by selling the content.

In either case, SPIES does not provide recompense to the content owner if content is released before time τ; it could be altered to provided such recompense, but that would give an incentive for the owner herself to register many times or discretely sell the content.

4 Verifying Content Possession

In the protocol, a hashing function H is used by registrants to prove possession of the content ϕ to the escrow service E without revealing ϕ directly. As long as ϕ is unchanged, a cryptographically secure hash such as MD5 is sufficient for proof of possession. This constitutes the simplest form of verification. However, either A or B may try to share ϕ without incurring negative financial consequences by first altering ϕ in some way. The recipient of the shared content would have content $\phi\prime$ with identical *semantic* content to ϕ, but with $H(\phi\prime) \neq H(\phi)$, and thus she could not use it to prove possession of ϕ to E. The ways in which ϕ could be perturbed without changing its meaning or function differ for different types of content; for video or audio content, changing encodings is sufficient, while other methods can be used for computer programs or text.

Enhancements to our protocol to defend against perturbation and sharing of ϕ also vary with the type of content. For content such as audio and video, the proof can be a hash of an extremely low quality version of the content. Details of the quality and encoding of the hash would be published by E at the beginning of the registration phase along with $d(\phi)$. Textual content can be converted into a sequence of hashes of sentences, where possession of some large proportion of these hashes constitutes proof of possession of ϕ (in fact, of text very closely related to ϕ). [2]. Computer programs can be subjected to static analysis, and this analysis, or a hash of it, be presented as proof of possession.

Finally, if a registrant believes she has a version of ϕ but is not able to prove possession, presumably because of successful perturbation by the sharer, she could call for manual verification during the Payment phase. A would submit ϕ to E, and the disputing registrant would submit her $\phi\prime$. A human acting as E would make a judgment on whether ϕ and $\phi\prime$ are semantically identical. All registrants, including A, would have committed to their version of ϕ during the Registration phase.

5 Conclusions

In this work, we present SPIES, a protocol that gives economic incentives to users of digital content to not distribute the content. With the SPIES protocol, digitial content providers can give out or sell their content in a limited release and expect rational participants to not put the content on a P2P file-sharing system for all to download. Even sharing amongst friends will be limited, as anyone with a copy of the content can get some payment, while the original user will lose a portion of her deposit. Users are also encouraged to be careful with the data, since a stolen copy can be used to obtain payment.

P2P systems will continue to have a negative connotation with many people as long as they are used primarily as a way to illegaly obtain copyrighted materials. We believe that, while DRM does not effectively stop widespread distribution for many content types, incentive-based schemes such as SPIES may provide a means to prevent the worst damages of this kind of abuse.

References

[1] The anonymizer. Available at http://www.anonymizer.com.

[2] Brin, S., Davis, J., and García-Molina, H. Copy Detection Mechanisms for Digital Documents. In *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data* (1995), pp. 398–409.

[3] Dingledine, R., Mathewson, N., and Syverson, P. Tor: The Second-Generation Onion Router. In *Proceedings of the 13th USENIX Security Symposium, forthcoming* (August 2004).

[4] Golle, P., Leyton-Brown, K., Mironov, I., and Lillibridge, M. Incentives for Sharing in Peer-to-Peer Networks. *Lecture Notes in Computer Science 2232* (2001).

[5] Horne, B., Pinkas, B., and Sander, T. Escrow Services and Incentives in Peer-to-Peer Networks. In *Proceedings of the 3rd ACM Conference on Electronic Commerce* (2001).

[6] Mathewson, N., and Dingledine, R. Mixminion: Strong Anonymity for Financial Cryptography. In *Proceedings of the Eighth International Conference on Financial Cryptography* (February 2004).

[7] Microsoft DRM Technologies Establish Foundation For Emerging Internet Music, Video and eBooks Industries. Microsoft Press Release, June 2001.

[8] Schneier, B. *Applied Cryptography.* John Wiley & Sons, 1996.

[9] Wayner, P. *Disappearing Cryptography, Second Edition – Information Hiding: Steganography and Watermarking.* Morgan Kaufmann, 2002.