

Traitor-Tracing on Binary Strings

Michael J. Collins

Sandia National Laboratories*
Albuquerque, NM USA 87185
mjcolli@sandia.gov

Abstract. Codes with the *Identifiable Parent Property* (IPP) have been studied in the context of traitor tracing; such codes can be used to enable a data supplier to determine the origin of pirated data. We consider an analogous property for a set of binary strings S : if a new string τ is formed by concatenating substrings of members of S , we should be able to identify at least one original string which must have been used to generate τ . We prove upper and lower bounds for the size of sets which satisfy this property.

Keywords: Traitor Tracing, Identifiable Parent Property, Strings, Watermarking

1 Introduction

Codes with the *Identifiable Parent Property* (IPP) were introduced in [3] (and generalized in [4]) with the motivation of detecting piracy when users combine several watermarked versions of a single document to produce a pirated version of the same document. We consider a related problem in which users generate new *derivative* documents by cutting and pasting from multiple watermarked documents. First we recall the definition of IPP codes:

Definition 1. Let \mathcal{C} be a code of length n over alphabet Σ , and let $T \subset \mathcal{C}$. Then $d = (d_1, d_2 \cdots d_n) \in \Sigma^n$ is a descendant of T and T is a parent set of d if, for each $1 \leq i \leq n$, there exists $(t_1, t_2 \cdots t_n) \in T$ with $d_i = t_i$.

If $d \in \Sigma^n$ has a parent set of size $\leq c$, then d is a c -descendant of \mathcal{C} ; a parent set of size $\leq c$ is a c -parent set.

Definition 2. Let \mathcal{C} be a code of length n over alphabet Σ . Then \mathcal{C} is c -IPP if, for any d which is a c -descendant of \mathcal{C} , the intersection of all c -parent sets of d is nonempty.

* Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

The idea is that the document D to be protected is divided into n segments, where each segment can be watermarked by embedding an element of Σ , and each legitimate document owner has a copy watermarked by a different code-word. Then c or fewer owners may collude to produce an unauthorized copy D' of D by combining portions of their copies. It is possible to identify at least one individual who must have contributed to D' by finding a parent of the vector of watermarks in D' [2]. This is a special case of the more general notion of *traitor tracing* [1].

1.1 c -IPP Sets of Strings

Denote concatenation of strings σ and τ by $\sigma\tau$. The length of a string σ is denoted by $|\sigma|$. Let σ^j denote the j th bit of σ , and define $\sigma[s : t] = \sigma^s\sigma^{s+1}\dots\sigma^t$. If s and t are not integers, then $\sigma[s : t] = \sigma[\lceil s \rceil : \lfloor t \rfloor]$.

A (sub)string of length k will be called a k -(sub)string. We write $\sigma \prec \tau$ to denote σ is a substring of τ . If s is a substring of some member of a set T , we write $s \prec T$.

We have a set U of “users”, with each user having a local copy of a collection of documents (i.e. binary strings) D . Associated to each user $i \in U$ there is a distinct string σ_i of length n . For each original document $d \in D$, user i has a watermarked version d_i of length m . The watermark string σ_i is embedded in d_i in such a way that, given a substring of d_i , we can extract the corresponding substring of σ_i ; i.e. there is an extraction function \mathcal{E} such that

$$\mathcal{E}(d_i[j : k]) = \sigma_i \left[j \frac{n}{m} : k \frac{n}{m} \right] .$$

Without loss of generality we can assume that all watermarked documents in D are of the same length m .

Now suppose one or more users create a new c -derivative document d' of length at least m by “cutting and pasting”, i.e. by concatenating at most c substrings of their various watermarked documents; we have

$$d' = d_{i_1}^{t_1}[j_1 : k_1] \cdots d_{i_c}^{t_c}[j_c : k_c]$$

where each $d_i^{t_i}$ is a copy of document d^t watermarked with σ_i . From d' we can then extract a string τ of length at least n which is the concatenation of the corresponding substrings of the σ_i , i.e.

$$\mathcal{E}(d') = \sigma_{i_1} \left[j_1 \frac{n}{m} : k_1 \frac{n}{m} \right] \cdots \sigma_{i_c} \left[j_c \frac{n}{m} : k_c \frac{n}{m} \right] .$$

Note that the extraction function \mathcal{E} must be able to extract this derivative watermark even though it is only given d' , without any indication of where the boundaries between its constituent substrings are.

Given a bound on c , we would like to use the extracted derivative watermark τ to identify (i.e. “trace”) at least one user who must have contributed to d' . Now we can ignore the outer documents and view users as simply combining substrings of their σ_i to produce τ . Thus we have the following definitions:

Definition 3. Let S be a set of binary strings of length n (i.e. $S \subset \{0, 1\}^n$), and let τ be a string of length n . Then $C \subset S$ is a c -parent set of τ if we can write $\tau = \tau_1\tau_2 \cdots \tau_c$ where each $\tau_i \prec C$. We say that τ is a c -descendant of C .

Note that we allow use of repeated and overlapping substrings.

Definition 4. The set $S \subset \{0, 1\}^n$ has the c -Identifiable Parent Property (c -IPP) if, for every c -descendant τ of S , the intersection of all c -parent sets of τ is nonempty.

Thus any member of this intersection can be identified as a parent of τ . The following definition is useful:

Definition 5. If $\sigma \in S$ and $\tau \prec \sigma$, then τ is unique with respect to S if it is not a substring of any other member of S . When S is clear from context, we just say that τ is unique.

2 Bounds for c -IPP Sets

2.1 Necessary Conditions

In order for S to be c -IPP, members of S must be sufficiently different from one another, which means that members of S must contain sufficiently short unique substrings. We have the following

Lemma 1. If S is c -IPP then there exists $\sigma \in S$ such that no $\lceil \frac{n}{1+\lfloor c/2 \rfloor} \rceil$ -substring of σ is a $\lfloor c/2 \rfloor$ -descendant of $S \setminus \{\sigma\}$.

Proof. Suppose on the contrary that every $\sigma \in S$ has a substring of length $\lceil \frac{n}{1+\lfloor c/2 \rfloor} \rceil$ which is a $\lfloor c/2 \rfloor$ -descendant of $S \setminus \{\sigma\}$. Then take $1 + \lfloor c/2 \rfloor$ such substrings from $1 + \lfloor c/2 \rfloor$ users and concatenate them to produce τ . None of these users can be identified as a parent of τ , since each one can be removed and replaced by the $\lfloor c/2 \rfloor$ other users who cover its contribution.

Now suppose σ is the string whose existence is guaranteed by lemma 1. Then it follows immediately that σ must have a unique substring of length

$$\left\lceil \frac{\lceil \frac{n}{1+\lfloor c/2 \rfloor} \rceil}{\lfloor c/2 \rfloor} \right\rceil = \left\lceil \frac{n}{\lfloor c^2/4 \rfloor + \lfloor c/2 \rfloor} \right\rceil$$

which gives

Theorem 1. If $S \subset \{0, 1\}^n$ is c -IPP then $|S| \leq 2^{\lceil n/(\lfloor c^2/4 \rfloor + \lfloor c/2 \rfloor) \rceil}$.

Proof. If we remove σ from S , the remaining set is still c -IPP, so by lemma 1 this smaller set also contains a unique substring of length $\lceil \frac{n}{\lfloor c^2/4 \rfloor + \lfloor c/2 \rfloor} \rceil$; if we repeatedly remove one string at a time, no unique substring can be removed twice.

For small values of c we can obtain better results than theorem 1:

Theorem 2. *If there are no unique $\lceil n/3 \rceil$ -substrings in S , then S is not 2-IPP; therefore if S is 2-IPP, $|S| \leq 2^{n/3}$.*

If there are no unique $\lceil n/5 \rceil$ -substrings in S then S is not 3-IPP; therefore if S is 3-IPP, $|S| \leq 2^{n/5}$.

Proof. Suppose there are no unique $\lceil n/3 \rceil$ -substrings in S . Take $X \in S$ and let σ be the middle third of X . Let $Y \neq X$ contain σ ; without loss of generality Y contains $\sigma\tau_Y$ where $|\tau_Y| \geq n/3$ and $\tau_Y \prec Z_Y \neq Y$ (see Fig. 1). Let τ_X be left third of X and let $\tau_X \prec Z_X \neq X$.

Now $\tau_X\sigma\tau_Y$ is a descendant of both $\{Z_X, Y\}$ and $\{Z_Y, X\}$; it may be the case that $Z_X = Z_Y$, but $\tau_X\sigma\tau_Y$ is also a descendant of $\{X, Y\}$. Therefore S is not 2-IPP.




Fig. 1. Upper bound for 2-IPP




Fig. 2. Upper bound for 3-IPP

Similarly, suppose there are no unique $\lceil n/5 \rceil$ -substrings in S . Given $X \in S$ let $\lambda_1\lambda_2\sigma \prec X$ with each λ_i, σ of length $\lceil n/5 \rceil$. By assumption there exists $Y \neq X$ with $\sigma\rho_1\rho_2 \prec Y$ with each ρ_i of length $\geq \lceil n/5 \rceil$ so that $|\lambda_1\lambda_2\sigma\rho_1\rho_2| = n$.

There also exist $X_1, X_2 \neq X$ such that $\lambda_i \prec X_i$ and $Y_1, Y_2 \neq Y$ such that $\rho_i \prec Y_i$. Thus $\lambda_1 \lambda_2 \sigma \rho_1 \rho_2$ is a descendant of $\{X, Y\}$ and of $\{X_1, X_2, Y\}$ and of $\{X, Y_1, Y_2\}$. Thus S is not 3-IPP (see Fig. 2).

2.2 Sufficient Conditions

Theorem 3. *If every $\lceil n/3 \rceil$ -substring of every member of S is unique, then S is 2-IPP. If every $\lceil n/5 \rceil$ -substring of every member of S is unique, then S is 3-IPP.*

Proof. Suppose S is not 2-IPP; we show that there must be a non-unique $\lceil n/3 \rceil$ -substring. Let s be a string which is a 2-descendant of S , but for which a parent cannot be identified. Then we can write $s = ab$ where a, b are substrings of $A, B \in S$. Without loss of generality let $|a| \geq \frac{n}{2}$. Since A cannot be identified as a parent, it must be possible to write $s = cd$ with $c \prec C \neq A$. If $|c| \geq \lceil n/3 \rceil$ then c is a non-unique substring shared by A and C . Otherwise, $|d| \geq \lceil 2n/3 \rceil$. Now $d \prec D \in S$, and (since D cannot be identified as a parent), we have $cd = ef$ with $e, f \prec E, F \neq D$: then at least one of E, F has a non-unique $\lceil n/3 \rceil$ -substring shared with D .

Similarly, suppose S is not 3-IPP; we show that there must be a non-unique $\lceil n/5 \rceil$ -substring. Let s be a 3-descendant of S for which a parent cannot be identified. As above, let $a, b, c, d \dots$ denote substrings of $A, B, C, D \dots \in S$. If we can write $s = abc$ with $\max(|a|, |b|, |c|) \geq \lceil 3n/5 \rceil$, then we immediately have a non-unique $\lceil n/5 \rceil$ -substring: align abc with a representation of s that does not use X , where $X \in \{A, B, C\}$ is the user contributing the longest of a, b, c . Then X 's contribution to s is split at no more than two points, yielding (at most) three non-unique substrings one of which must have length at least $\lceil n/5 \rceil$.

Otherwise let $s = abc$ and (with no loss of generality) let $|a| \geq |c|$. Now suppose $\lceil 2n/5 \rceil \leq |a| < \lceil 3n/5 \rceil$, and let $abc = def$ with $A \neq D, E, F$. If $\max(|d|, |e|) \geq \lceil n/5 \rceil$, then either D or E has a non-unique $\lceil n/5 \rceil$ -substring shared with A (see Fig. 3). Otherwise $|f| \geq \lceil 3n/5 \rceil$.

Now suppose $\lceil n/5 \rceil \leq |a| < \lceil 2n/5 \rceil$, and let $abc = def$ with $A \neq D, E, F$. If $|d| \geq \lceil n/5 \rceil$ then D has a non-unique $\lceil n/5 \rceil$ -substring shared with A . Otherwise $\max(|e|, |f|) \geq \lceil 2n/5 \rceil$. If $|f| \geq \lceil 2n/5 \rceil$ we have essentially the same situation as Fig. 3 (with f in place of a).

So suppose $|e| \geq \lceil 2n/5 \rceil, f < \lceil 2n/5 \rceil, d < \lceil n/5 \rceil$ and let $def = ghi$ with $E \neq G, H, I$. If $\max(|g|, |i|) \geq \lceil 2n/5 \rceil$ we are again in the same situation as Fig. 3 (with g or i in place of a). Otherwise we have $|h| \geq \lceil n/5 \rceil$. Now this implies that E must have a non-unique substring; either $h \prec e$, or the alignment of def with ghi splits e at only one point.

Finally, if $|a| < \lceil n/5 \rceil$ then $|b| \geq \lceil 3n/5 \rceil$ (since $|a| \geq |c|$).

3 c-Traceability

The identifiable parent property only guarantees that a parent of the descendant string τ can be identified if all possible c -parent sets are considered. We would




Fig. 3. Shaded region is non-unique substring of length $\geq \lceil n/5 \rceil$

like to have an efficient algorithm for actually identifying at least one of the parents whose existence is guaranteed. It is natural to consider the $\sigma \in S$ which has the longest common substring with τ , but such a σ need not be a parent of τ in general. Given strings a, b let $\ell(a, b)$ denote the length of a longest common substring.

Definition 6. Let $S \subset \{0, 1\}^n$ be c -IPP. S has the c -traceability property (c -TA) if, for every c -descendant τ of S , the intersection of all c -parent sets of τ includes all $\sigma \in S$ which maximize $\ell(\sigma, \tau)$.

We have the following

Theorem 4. If every substring in $S \subset \{0, 1\}^n$ of length

$$\frac{4n}{(c+3)^2}$$

is unique, then S is c -TA.

Proof. Let τ be a c -descendant of S , so $\tau = s_1 s_2 \cdots s_c$ where $|\tau| = n$ and each s_i is a substring of a member of $T \subset S$. Suppose S is not c -TA, so there exists $r \in S \setminus T$ such that $\ell(r, \tau) \geq |s_i|$ for all i . Let λ be such a longest common substring with $|\lambda| = \frac{n}{\theta}$ and $\lambda \prec s_i s_{i+1} \cdots s_{i+p}$. Then (see Fig. 4)

$$|s_1 \cdots s_i s_{i+p} \cdots s_c| \geq n \frac{\theta - 1}{\theta}$$

which implies there is some s_j such that

$$|s_j| \geq \frac{n \frac{\theta - 1}{\theta}}{c - p + 1}.$$

But $|s_j| \leq \frac{n}{\theta}$ which requires

$$p \leq c + 2 - \theta.$$

Now one of s_i, \dots, s_{i+p} must contain a non-unique substring of length at least $\frac{n}{\theta(p+1)}$, i.e. at least

$$\frac{n}{\theta(c+3-\theta)}.$$

The denominator is maximized at $\theta = (c + 3)/2$, giving a non-unique substring of length at least

$$\frac{4n}{(c + 3)^2} .$$




Fig. 4. Proof of Theorem 4

4 Inclusion of Extraneous Substrings

We now consider the possibility that, when users combine portions of watermarked documents, some portions of non-watermarked documents may be included. If the amount of such *extraneous* content is limited, we can still trace parents. We will assume that the watermark extraction algorithm \mathcal{E} , when applied to unmarked input, produces essentially random results; thus any sequence of bits can appear in the extraneous portion of the watermark.

Suppose that the total length of the extraneous watermark is $\leq \alpha n$. A derivative watermark τ is produced by concatenation of no more than c substrings, some of which now may be extraneous. Now a (c, α) -descendant of S is a string $\tau = s_1 \cdots s_c$ with $|\tau| = n$ such that

$$\sum_{i:s_i \not\in S} |s_i| \leq \alpha n$$

We say that S is (c, α) -TA if it is still the case that any σ_i maximizing $\ell(\sigma_i, t)$ must be a parent of t .

Theorem 5. *Let $\alpha < \frac{1}{c+3}$. Then $S \subset \{0, 1\}^n$ is (c, α) -TA if all substrings of length*

$$\left(\frac{1 + \beta}{c + 3} \right)^2 n$$

are unique, where $\beta = \sqrt{1 - \alpha(c + 3)}$

Proof. Using the same notation and argument as in the proof of theorem 4, we now have that the non-extraneous members of the set $\{s_i, \dots, s_{i+p}\}$ must include a non-unique substring length at least $\frac{n/\theta - \alpha n}{p+1}$, i.e. at least

$$\frac{n(1 - \alpha\theta)}{\theta(c + 3 - \theta)} .$$

This is minimized at

$$\theta = \frac{1 - \sqrt{1 - \alpha(c + 3)}}{\alpha}.$$

We have defined β such that $\alpha = \frac{1 - \beta^2}{c + 3}$. Then the lower bound is minimized at $\theta = \frac{c + 3}{1 + \beta}$, so we have a non-unique substring of length at least

$$\left(\frac{1 + \beta}{c + 3}\right)^2 n.$$

5 Lower Bound

For any $k < n$ we can construct a set of strings of length n in which every k -substring is unique as follows. Divide $\{0, 1\}^k$ into equivalence classes, where the equivalence class of $r = r_1 r_2 \cdots r_k$ consists of all rotations of r , i.e. all strings of the form $r_t r_{t+1} \cdots r_k r_1 \cdots r_{t-1}$. There are at least $\lceil 2^k/k \rceil$ such classes. Select a representative r_i from each class, and form σ_i by concatenating $\lfloor n/k \rfloor$ copies of r_i , followed by the first $n - k\lfloor n/k \rfloor$ bits of r_i to get $|\sigma_i| = n$. Thus we have

Theorem 6. *For any c, n there exists $S \subset \{0, 1\}^n$ which is c -TA such that*

$$|S| \geq \frac{2^{\lceil 4n/(c+3)^2 \rceil}}{\lceil 4n/(c+3)^2 \rceil}.$$

For any n there exists $S \subset \{0, 1\}^n$ which is 2-IPP such that

$$|S| \geq \frac{2^{\lceil n/3 \rceil}}{\lceil n/3 \rceil}.$$

For any n there exists $S \subset \{0, 1\}^n$ which is 3-IPP such that

$$|S| \geq \frac{2^{\lceil n/5 \rceil}}{\lceil n/5 \rceil}.$$

References

1. Benny Chor, Amos Fiat, Moni Naor, and Benny Pinkas. Tracing traitors. *IEEE Transactions on Information Theory*, 46(3):893–910, 2000.
2. Amos Fiat and Tamir Tassa. Dynamic traitor tracing. *Proceedings of Crypto '99, LNCS*, 1666:354–371, 1999.
3. Henk D. L. Hollmann, Jack H. van Lint, Jean-Paul Linnartz, and Ludo M. G. M. Tolhuizen. On codes with the identifiable parent property. *J. Combinatorial Theory, Series A*, 82(2):121–133, 1998.
4. Jessica Staddon, Douglas R. Stinson, and Ruizhong Wei. Combinatorial properties of frameproof and traceability codes. *IEEE Transactions on Information Theory*, 47(3):1042–1049, 2001.