Visual Cryptography*

Moni Naor † Adi Shamir ‡

Abstract

In this paper we consider a new type of cryptographic scheme, which can decode concealed images without any cryptographic computations. The scheme is perfectly secure and very easy to implement. We extend it into a visual variant of the k out of n secret sharing problem, in which a dealer provides a transparency to each one of the n users; any k of them can see the image by stacking their transparencies, but any $k-1$ of them gain no information about it.

*A preliminary version of this paper appeared in Eurocrypt 94.

†Dept. of Applied Mathematics and Computer Science, Weizmann Institute of Science, Rehovot 76100, Israel. E-mail: naor@wisdom.weizmann.ac.il. Research supported by an Alon Fellowship and a grant from the Israel Science Foundation administered by the Israeli Academy of Sciences.

‡Dept. of Applied Mathematics and Computer Science, Weizmann Institute of Science, Rehovot 76100, Israel. E-mail: shamir@wisdom.weizmann.ac.il.

1 Introduction

In this paper we consider the problem of encrypting written material (printed text, handwritten notes, pictures, etc.) in a perfectly secure way which can be decoded directly by the human visual system. The basic model consists of a printed page of ciphertext (which can be sent by mail or faxed) and a printed transparency (which serves as a secret key). The original cleartext is revealed by placing the transparency with the key over the page with the ciphertext, even though each one of them is indistinguishable from random noise. The system is similar to a one time pad in the sense that each page of ciphertext is decrypted with a different transparency. Due to its simplicity, the system can be used by anyone without any knowledge of cryptography and without performing any cryptographic computations.

The best way to visualize the visual cryptographic scheme is to consider a concrete example. At the end of the paper we enclose two random looking dot patterns. To decrypt the secret message, the reader should photocopy each pattern on a separate transparency, align them carefully, and project the result with an overhead projector.

This basic model can be extended into a visual variant of the k out of n secret sharing problem: Given a written message, we would like to generate n transparencies so that the original message is visible if any k (or more) of them are stacked together, but totally invisible if fewer than k transparencies are stacked together (or analysed by any other method). The original encryption problem can be considered as a 2 out of 2 secret sharing problem.

The main results of this paper (besides introducing this new paradigm of cryptographic schemes) include practical implementations of a k out of n visual secret sharing scheme for small values of k and n, as well as efficient asymptotic constructions which can be proven optimal within certain classes of schemes.

2 The Model

The simplest version of the visual secret sharing problem assumes that the message consists of a collection of black and white pixels and each pixel is handled separately[1]. Each original pixel appears in n modified versions (called shares), one for each transparency. Each share is a collection of m black and white subpixels, which are printed in close proximity to each other so that the human visual system averages their individual black/white contributions. The resulting structure can be described by an $n \times m$ Boolean matrix $S = [s_{ij}]$ where $s_{ij} = 1$ iff the jth subpixel in the ith transparency is black. When transparencies $i_1, i_2, \ldots i_r$ are stacked together in a way which properly aligns the subpixels, we see a combined share whose black subpixels are represented by the Boolean "or" of rows $i_1, i_2, \ldots i_r$ in S. The grey level of this combined share is proportional to the Hamming weight $H(V)$ of the "or"ed m-vector V. This grey level is interpreted by the visual system of the users as black if $H(V) \geq d$ and as white if $H(V) < d - \alpha m$ for some fixed threshold $1 \leq d \leq m$ and relative difference $\alpha > 0$.

This framework resembles the framework of linear codes, with the important difference

[1] It is conceivable that handling larger groups of pixels simultaneously yields better results

that the underlying algebraic structure is a semi-group rather than a group. In particular, the visual effect of a black subpixel in one of the transparencies cannot be undone by the colour of that subpixel in other transparencies which are laid over it. This monotonicity rules out common encryption techniques which add random noise to the cleartext during the encryption process, and subtracts the same noise from the ciphertext during the decryption process. It also rules out the more natural model in which a white pixel is represented by a completely white collection of subpixels and a black pixel is represented by a completely black collection of subpixels, and thus we have to use a threshold d and relative difference $\alpha > 0$ to distinguish between the colours.

Definition 2.1 *A solution to the k out of n visual secret sharing scheme consists of two collections of $n \times m$ Boolean matrices $\mathcal{C}_0$ and $\mathcal{C}_1$. To share a white pixel, the dealer randomly chooses one of the matrices in $\mathcal{C}_0$, and to share a black pixel, the dealer randomly chooses one of the matrices in $\mathcal{C}_1$. The chosen matrix defines the colour of the m subpixels in each one of the n transparencies. The solution is considered valid if the following three conditions are met:*

1. *For any S in $\mathcal{C}_0$, the "or" V of any k of the n rows satisfies $H(V) \leq d - \alpha \cdot m$.*
2. *For any S in $\mathcal{C}_1$, the "or" V of any k of the n rows satisfies $H(V) \geq d$.*
3. *For any subset $\{i_1, i_2, \ldots i_q\}$ of $\{1, 2, \ldots n\}$ with $q < k$, the two collections of $q \times m$ matrices $\mathcal{D}_t$ for $t \in \{0, 1\}$ obtained by restricting each $n \times m$ matrix in C_t (where $t = 0, 1$) to rows $i_1, i_2, ..., i_q$ are indistinguishable in the sense that they contain the same matrices with the same frequencies.*

Condition 3 implies that by inspecting fewer than k shares, even an infinitely powerful cryptanalyst cannot gain any advantage in deciding whether the shared pixel was white or black. In most of our constructions, there is a function f such that the combined shares from $q < k$ transparencies consist of all the V's with $H(V) = f(q)$ with uniform probability distribution, regardless of whether the matrices were taken from $\mathcal{C}_0$ or $\mathcal{C}_1$. Such a scheme is called *uniform*. The first two conditions are called *contrast* and the third condition is called *security*.

The important parameters of a scheme are:

- m, the number of pixels in a share. This represents the loss in resolution from the original picture to the shared one. We would like m to be as small as possible.
- α, the relative difference in weight between combined shares that come from a white pixel and a black pixel in the original picture. This represents the loss in contrast. We would like α to be as large as possible.
- r, the size of the collections $\mathcal{C}_0$ and $\mathcal{C}_1$ (they need not be the same size, but in all of our constructions they are). $\log r$ represents the number of random bits needed to generate the shares and does not effect the quality of the picture.

Results: We have a number of constructions for specific values of k and n. For general k we have a construction for the k out k problem with $m = 2^{k-1}$ and $\alpha = \frac{1}{2^{k-1}}$ and we have a proof of optimality of this scheme. For general k and n we have a construction with $m = \log n \cdot 2^{O(k \log k)}$ and $\alpha = \frac{1}{2^{\Omega(k)}}$.

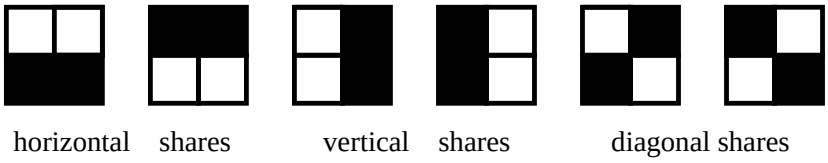

Figure 1:

3 Efficient solutions for small k and n

The 2 out of n visual secret sharing problem can be solved by the following collections of $n \times n$ matrices:

$$\mathcal{C}_0 = \{\text{all the matrices obtained by permuting the columns of } \begin{bmatrix} 100\ldots0 \\ 100\ldots0 \\ \ldots \\ 100\ldots0 \end{bmatrix}\}$$

$$\mathcal{C}_1 = \{\text{all the matrices obtained by permuting the columns of } \begin{bmatrix} 100\ldots0 \\ 010\ldots0 \\ \ldots \\ 000\ldots1 \end{bmatrix}\}$$

Any single share in either $\mathcal{C}_0$ or $\mathcal{C}_1$ is a random choice of one black and $n-1$ white subpixels. Any two shares of a white pixel have a combined Hamming weight of 1, whereas any two shares of a 1 pixel have a combined Hamming weight of 2, which looks darker. The visual difference between the two cases becomes clearer as we stack additional transparencies.

The original problem of visual cryptography is the special case of a 2 out of 2 visual secret sharing problem. It can be solved with two subpixels per pixel, but in practice this can distort the aspect ratio of the original image. It is thus recommended to use 4 subpixels arranged in a 2×2 array where each share has one of the visual forms in Figure 1. A white pixel is shared into two identical arrays from this list, and a black pixel is shared into two complementary arrays from this list. Any single share is a random choice of two black and two white subpixels, which looks medium grey. When two shares are stacked together, the result is either medium grey (which represents white) or completely black (which represents black).

The next case is the 3 out of 3 visual secret sharing problem, which is solved by the following scheme:

$$\mathcal{C}_0 = \{\text{all the matrices obtained by permuting the columns of } \begin{bmatrix} 0011 \\ 0101 \\ 0110 \end{bmatrix}\}$$

$$\mathcal{C}_1 = \{\text{all the matrices obtained by permuting the columns of } \begin{bmatrix} 1100 \\ 1010 \\ 1001 \end{bmatrix}\}$$

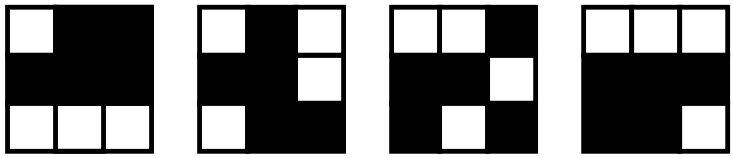

shares of a white pixel

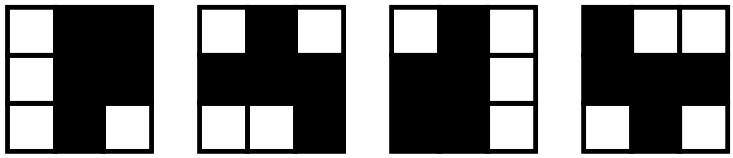

shares of a black pixel

Figure 2:

Note that the six shares described by the rows of $\mathcal{C}_0$ and $\mathcal{C}_1$ are exactly the six 2×2 arrays of subpixels from Fig. 1. Each matrix in either $\mathcal{C}_0$ or $\mathcal{C}_1$ contains one horizontal share, one vertical share and one diagonal share. Each share contains a random selection of two black subpixels, and any pair of shares from one of the matrices contains a random selection of one common black subpixel and two individual black subpixels. Consequently, the analysis of one or two shares makes it impossible to distinguish between $\mathcal{C}_0$ and $\mathcal{C}_1$. However, a stack of three transparencies from $\mathcal{C}_0$ is only 3/4 black, whereas a stack of three transparencies from $\mathcal{C}_1$ is completely black.

The following scheme generalizes this 3 out of 3 scheme into a 3 out of n scheme for an arbitrary $n \geq 3$. Let B be the black $n \times (n-2)$ matrix which contains only 1's, and let I be the identity $n \times n$ matrix which contains 1's on the diagonal and 0's elsewhere. Let BI denote the $n \times (2n-2)$ matrix obtained by concatenating B and I, and let $c(BI)$ be the Boolean complement of the matrix BI. Then

$\mathcal{C}_0 = \{$all the matrices obtained by permuting the columns of $c(BI)\}$

$\mathcal{C}_1 = \{$all the matrices obtained by permuting the columns of $BI\}$

has the following properties: Any single share contains an arbitrary collection of $n-1$ black and $n-1$ white subpixels; any pair of shares have $n-2$ common black and two individual black subpixels; any stacked triplet of shares from $\mathcal{C}_0$ has n black subpixels, whereas any stacked triplet of shares from $\mathcal{C}_1$ has $n+1$ black subpixels.

The 4 out of 4 visual secret sharing problem can be solved by the shares described in Figure 2 (along with all their permutations).

Any single share contains 5 black subpixels, any stacked pair of shares contains 7 black subpixels, any stacked triplet of shares contains 8 black subpixels, and any stacked quadruple of shares contains either 8 or 9 black subpixels, depending on whether the shares were taken from $\mathcal{C}_0$ or $\mathcal{C}_1$. It is possible to reduce the number of subpixels from 9 to 8, but then they cannot be packed into a square array without distorting their aspect ratio.

Finally, we describe an efficient 2 out of 6 scheme. The scheme is defined by

$$\mathcal{C}_0 = \{ \text{ all the matrices obtained by permuting the columns of } \begin{bmatrix} 1100 \\ 1100 \\ 1100 \\ 1100 \\ 1100 \\ 1100 \end{bmatrix} \}$$

$$\mathcal{C}_1 = \{\text{all the matrices obtained by permuting the columns of } \begin{bmatrix} 0101 \\ 1010 \\ 1100 \\ 0011 \\ 0110 \\ 1001 \end{bmatrix} \}$$

The scheme has contrast $\frac{1}{4}$: any two shares of $\mathcal{C}_0$ cover 2 out of 4 of the pixels, while any pair of shares from $\mathcal{C}_1$ covers at least 3 out of 4 pixels (some cover all four). The security of the scheme follows from the fact that in both $\mathcal{C}_0$ and $\mathcal{C}_1$ each share is random subset of 2 black pixels out of 4.

One possible generalization of this scheme to a 2 out of n scheme is to fix m so that $\binom{m}{m/2} \geq n$ and consider all subsets of size $m/2$ of some ground set of size m. The ith row is S^1 corresponds to the ith subset, i.e. $S^1[i,j] = 1$ iff the jth element is in the ith subset. S^0 is the $n \times m$ matrix where each row is $1^{m/2}0^{m/2}$. $\mathcal{C}_0$ and $\mathcal{C}_1$ are obtained from all column permutations of S^0 and S^1. The contrast achieved this way is $1/m$. As we shall see in Section 5, we can do better than that.

4 A general k out of k scheme

We now describe two general constructions which can solve any k out of k visual secret sharing problem by using 2^k and 2^{k-1} subpixels respectively. We then prove that the second construction is optimal in that any k out k scheme must use at least 2^{k-1} pixels.

4.1 Construction 1

To define the two collections of matrices we make use of two lists of vectors $J_1^0, J_2^0, \ldots J_k^0$ and $J_1^1, J_2^1, \ldots J_k^1$. Let $J_1^0, J_2^0, \ldots J_k^0$ be vectors of length k over $GF[2]$ with the property that every $k-1$ of them are linearly independent over $GF[2]$, but the set of all k vectors is not independent. Such a collection can be easily constructed, e.g. let $J_i^0 = 0^{i-1}10^{k-i}$ for $1 \leq i \leq k$ and $J_k^0 = 1^{k-1}0$. Let $J_1^1, J_2^1, \ldots J_k^1$ be vectors of length k over $GF[2]$ with the property that they are linearly independent over $GF[2]$. (This can be thought of as a first order Reed-Muller code [7])

Each list defines a $k \times 2^k$ matrix S^t for $t \in \{0,1\}$ and the collections $\mathcal{C}_0$ and $\mathcal{C}_1$ are obtained by permuting the columns of the corresponding matrix in all possible ways. We index the columns of S^t by vectors of length k over $GF[2]$. For $t \in \{0,1\}$ let S^t be defined

as follows: $S^t[i,x] \doteq < J_i^t, x >$ for any $1 \leq i \leq k$ and any vector x of length k over $GF[2]$ where $< x, y >$ denotes the inner product over $GF[2]$.

Lemma 4.1 *The above scheme is a k out of k scheme with parameters $m = 2^k$, $\alpha = 1/2^k$ and $r = 2^k!$.*

Proof: In order to show contrast, note that in matrix S^0 there are two columns that are all zero; in the example given these are the column indexed by $\vec{x} = 0^k$ and the column indexed by $\vec{x} = 0^{k-1}1$. On the other hand, in S^1 there is only one column that is all 0, the one corresponding to $\vec{x} = 0^k$. Therefore in any permutation of S^0 the "or" of the k rows yields $2^k - 2$ ones, whereas in any permutation of S^1 the "or" of the k rows yields $2^k - 1$ ones.

In order to show security, note that the vectors corresponding to any $k-1$ rows in both S^0 and S^1 are linearly independent over $GF[2]$. Therefore if one considers the rows as subsets of a ground set of size 2^k, then every intersection of $k-1$ rows or their complement has the same size, two. (Note that we include complemented sets, and thus if all possible intersections of $k-1$ are the same, then all smaller intersections are the same as well.) In other words, consider the columns in S^0 and S^1 obtained by restricting to the $k-1$ chosen rows. Then every possible assignment to the $k-1$ entries appears exactly twice. Hence, a random permutation of the columns, as is used to generate $\mathcal{C}_0$ and $\mathcal{C}_1$, yields the same distribution regardless of which $k-1$ rows were chosen. □

4.2 Construction 2

We now show a slightly better scheme with parameters $m = 2^{k-1}$, $\alpha = 1/2^{k-1}$ and $r = 2^{k-1}!$. Consider a ground set $W = \{e_1, e_2, \ldots e_k\}$ of k elements and let $\pi_1, \pi_2, \ldots \pi_{2^{k-1}}$ be a list of all the subsets of even cardinality and let $\sigma_1, \sigma_2, \ldots \sigma_{2^{k-1}}$ be a list of all the subsets of W of odd cardinality (the order is not important).

Each list defines the following $k \times 2^{k-1}$ matrices S^0 and S^1: For $1 \leq i \leq k$ and $1 \leq j \leq 2^{k-1}$ let $S^0[i,j] = 1$ iff $e_i \in \pi_j$ and $S^1[i,j] = 1$ iff $e_i \in \sigma_j$.

As in the construction above, the collections $\mathcal{C}_0$ and $\mathcal{C}_1$ are obtained by permuting all the columns of the corresponding matrix.

Lemma 4.2 *The above scheme is a k out of k scheme with parameters $m = 2^{k-1}$, $\alpha = 1/2^{k-1}$ and $r = 2^{k-1}!$.*

Proof: In order to show contrast, note the in matrix S^0 there is one column that is all zero, the one indexed by the empty set. On the other hand, in S^1 there is no column that is all 0. Therefore in any permutation of S^0 the "or" of the k rows yields only $2^{k-1} - 1$ ones, whereas in any permutation of S^1 the "or" of the k rows yields 2^{k-1} ones.

In order to show security, note that if one examines any $k-1$ rows in either S^0 and S^1 then the structure discovered is similar: consider the rows as subsets of a ground set of size 2^{k-1}; every intersection of $k-1$ rows or their complement has the same size, two. Hence, as in the proof of Lemma 4.1, a random permutation of the columns yields the same distribution regardless of which $k-1$ rows were chosen. □

4.3 Upper bound on α

We show that α must be exponentially small as a function of k and, in fact, get a tight bound that $\alpha \geq 2^{k-1}$. The key combinatorial fact used is the following (see [5, 6]: given two sequences of sets $A_1, A_2, \ldots A_k$ and $B_1, B_2, \ldots B_k$ of some ground set G such that for every subset $U \subset \{1, ..k\}$ of size at most $k-1$ we have $|\bigcap_{i \in U} A_i| = |\bigcap_{i \in U} B_i|$, then $|\cup_{i=1}^{k} A_i| \leq \frac{1}{2^{k-1}} \cdot |G| + |\cup_{1=1}^{k} B_i|$. In other words, if the intersections of the A_i's and B_i's agree in size for all subsets smaller than k elements, then the difference in the union cannot be too large.

Consider now a k out k scheme $\mathcal{C}$ with parameters m, α and r. Let the two collections be $\mathcal{C}_0$ and $\mathcal{C}_1$. We construct from the collections two sequences of sets $A_1, A_2, \ldots A_k$ and $B_1, B_2, \ldots B_k$. The ground set is of size $m \cdot r$ and its elements are indexed by (x, y) where $1 \leq x \leq r$ and $1 \leq y \leq m$. Element (x, y) is in A_i iff $S_x^0[iy] = 1$ and element (x, y) is in B_i iff $S_x^1[iy] = 1$.

We claim that for any $U \subset \{1, ..k\}$ of size $q < k$ the equality $|\bigcap_{i \in U} A_i| = |\bigcap_{i \in U} B_i|$ holds: the security condition of $\mathcal{C}$ implies that we can construct a 1-1 mapping between all the $q \times m$ matrices obtained from considering only rows corresponding to U in $\mathcal{C}_0$ and the $q \times m$ matrices of $\mathcal{C}_1$ such that any two matched matrices are identical. (Strictly speaking, the security condition is not strong enough to imply it, but given any scheme we can convert it into one that has this property without changing α and m.) Therefore when considering $|\bigcap_{i \in U} A_i|$ and $|\bigcap_{i \in U} B_i|$ the contribution of each member of a pair of matched matrices is identical and hence $|\bigcap_{i \in U} A_i| = |\bigcap_{i \in U} B_i|$.

Applying now the combinatorial fact mentioned above yields that

$$|\cup_{i=1}^{k} B_i| \leq \frac{1}{2^{k-1}} \cdot rm + |\cup_{1=1}^{k} A_i|.$$

This means that for at least one matrix in $\mathcal{C}_1$ and one matrix in $\mathcal{C}_0$ the difference between the Hamming weight of the "or" of their rows is at most $\frac{1}{2^{k-1}} \cdot m$. Hence we have:

Theorem 4.3 *In any k out k scheme $\alpha \leq \frac{1}{2^{k-1}}$ and $m \geq 2^{k-1}$.*

5 A general k out of n scheme

In this section we construct a k out of n scheme. What we show is how to go from a k out of k scheme to a k out of n scheme.

Let $\mathcal{C}$ be an k out of k visual secret sharing scheme with parameters m, r, α. The scheme $\mathcal{C}$ consists of two collections of $k \times m$ Boolean matrices $\mathcal{C}_0 = T_1^0, T_2^0, \ldots T_r^0$ and $\mathcal{C}_1 = T_1^1, T_2^1, \ldots T_r^1$. Furthermore, assume the scheme is uniform, i.e. there is a function $f(q)$ such that for any matrix T_i^t where $t \in \{0, 1\}$ and $1 \leq i \leq r$ and for every $1 \leq q \leq k-1$ rows of T_i^t the Hamming weight of the "or" of the q rows is $f(q)$. Note that all our previous constructions have this property.

Let H be a collection of ℓ functions such that

1. $\forall h \in H$ we have $h : \{1..n\} \mapsto \{1..k\}$
2. For all subsets $B \subset \{1..n\}$ of size k and for all $1 \leq q \leq k$ the probability that a randomly chosen $h \in H$ yields q different values on B is the same. Denote this probability by β_q

We construct from $\mathcal{C}$ and H a k out of n scheme $\mathcal{C}'$ as follows:

- The ground set is $V = U \times H$ (i.e. it is of size $m \cdot \ell$ and we consider its elements as indexed by a member of U and a member of H).
- Each $1 \leq t \leq r^\ell$ is indexed by a vector $(t_1, t_2, \ldots t_\ell)$ where each $1 \leq t_i \leq r$.
- The matrix S_t^b for $t = (t_1, t_2, \ldots t_\ell))$ where $b \in \{0, 1\}$ is defined as

$$S_t^b[i, (j, h)] = T_{t_h}^b[h(i), j]$$

Note that in the above expression t_h means the hth entry in t, where h is simply interpreted as a number between 1 and ℓ.

Lemma 5.1 *If $\mathcal{C}$ is a scheme with parameters m, α, r, then $\mathcal{C}'$ is a scheme with parameters $m' = m \cdot \ell, \alpha' = \alpha \cdot \beta_k, r' = r^\ell$.*

Proof: In order to show contrast, note that for any k rows in a matrix S_t^b and any $h \in H$, if the subset corresponding to the k rows is mapped to $q < k$ different values by h, then we know by the assumption of uniformity that the weight of the "or" of the q rows in $\mathcal{C}$ is $f(q)$. The difference between white pixels and black pixels occurs only when h is $1-1$ which happens at β_k of the $h \in H$ and it is $\alpha \cdot m$ in this case. Therefore the Hamming weight of an "or" of k rows of a white pixel is at most

$$\ell(\beta_k \cdot (d - \alpha m) + \sum_{q=1}^{k-1} \beta_q \cdot f(q))$$

and the weight of a black pixel is

$$\ell(\beta_k \cdot d + \sum_{q=1}^{k-1} \beta_q \cdot f(q))$$

which means that the relative difference between them is at least $\beta_k \cdot \alpha$.

In order to see the security of the scheme, note that we are essentially repeating ℓ times the scheme $\mathcal{C}$ where each instance is independent of all other instances. Therefore from the security of $\mathcal{C}$ we get the security of $\mathcal{S}$. □

5.1 Construction of H

One can construct H from a collection of k-wise independent hash functions (see e.g. [3], [4], [9]). Suppose that H is such that for any k values $x_1, x_2, \ldots x_k \in \{1, ..n\}$ the k random variables defined by $X_1 \doteq h(x_1), X_2 \doteq h(x_2), \ldots X_k \doteq h(x_k)$ for a randomly chosen $h \in H$ are completely independent. Since they are independent, the probability that they yield q different values is the same, no matter what $x_1, x_2, \ldots x_k$ are. For a concrete example, assume that k is a prime (otherwise we have to deal with its factors), and let l be such that $k^l \geq n$. The family H is based on the set of polynomials of degree $k-1$ over $GF[k^l]$, where

for ever $h \in H$ there is a corresponding polynomial q(x), and $h(x) = q(x) \bmod k$. The size of H is about n^k. The probability β_k that a random h is $1-1$ on a set of k elements is

$$\frac{k!}{k^k} \geq \frac{(k/e)^k}{k^k\sqrt{2\pi k}} = \frac{e^{-k}}{\sqrt{2\pi k}}.$$

We can therefore conclude by applying Lemma 5.1:

Theorem 5.2 *For any n and k there exists a visual secret sharing scheme with parameters $m = n^k \cdot 2^{k-1}$, $\alpha = (2e)^{-k}/\sqrt{2\pi k}$ and $r = n^k(2^{k-1}!)$.*

5.2 Relaxing the conditions on H

Suppose now that we relax Condition 2 in the definition of H to the following: there exists an ϵ such that for all subsets $B \subset \{1..n\}$ of size k and for all $1 \leq q \leq k$ the probability that a randomly chosen $h \in H$ yields q different values on B is the same to within ϵ. As we shall see, this leeway allows for much smaller H's.

Taking ϵ to be small, say smaller than $\alpha\beta_k/4$, cannot make a big difference in the quality of our construction: The Hamming weight of an "or" of k rows of a white pixel is at most

$$\ell((\beta_k + \epsilon) \cdot (d - \alpha m) + \sum_{q=1}^{k-1} (\beta_q + \epsilon) \cdot f(q))$$

and the weight of a black pixel is at least

$$\ell((1-\epsilon)\beta_k \cdot d + \sum_{q=1}^{k-1} (1-\epsilon) \cdot \beta_q \cdot f(q)).$$

The relative difference between black and white is therefore at least $\beta_k \cdot \alpha - 2\epsilon$.

Note that the security of the scheme is not effected at all, since fewer than k shares never map to k different values.

Construction of relaxed H:

We use *small-bias probability spaces* to construct such a relaxed family (see [8], [2], [3] for definitions and constructions). A probability space with random variables that are ϵ-bias is an approximation to a probability space with completely independent random variables, in that the bias (i.e. the difference between the probability that there parity is 0 and 1) is bounded by ϵ (as opposed to 0 in the complete independence). Similarly, a probability space which is k-wise ϵ-bias is an approximation to k-wise independent probability spaces.

Assume that k is a power of 2. Let R be a $k \log k$-wise δ-bias probability space on $n \log k$ random variables which takes values in $\{0,1\}$. They are indexed as Y_{ij} for $1 \leq i \leq n$ and $1 \leq j \leq \log k$. There are explicit constructions of such probability spaces of size $2^{O(k \log k)} \log n$ (see [8] [1]).

Each function h corresponds to a point in the probability space. $h(x)$ is the value of $Y_{x1}, Y_{x2}, \ldots, Y_{x \log k}$ treated as a number between 0 and $2^k - 1$. It can be shown that for all $x_1, x_2, \ldots x_k \in \{1, ..n\}$ and for all $y_1, y_2, \ldots y_k \in \{0, ..2^k - 1\}$ we have

$$\frac{1}{k^k} - \delta \cdot k^k \leq Prob[h(x_1) = y_1, h(x_2) = y_2, \ldots h(x_k) = y_k] \leq \frac{1}{k^k} + \delta k^k.$$

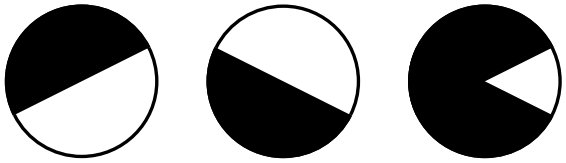

Figure 3:

Therefore taking $\delta = \frac{1}{4k^{2k}}$ implies that $\epsilon \leq 2^{-2k}$ and we get a scheme in which the number of subpixels grows only logarithmically with the number of shares n.

Theorem 5.3 *For any n and k there exists a visual secret sharing scheme with parameters $m = \log n \cdot 2^{O(k \log k)}$, $\alpha = 2^{-\Omega(k)}$.*

We do not know whether the bound on m in the above theorem is tight, but we conjecture that $\log n 2^{O(k)}$ is the right answer.

6 Extensions

There are many possible enhancements and extensions of the basic model introduced in this paper. Consider, for example, the problem of visual encryption of a continuous tone image whose pixels have grey levels ranging from 0 to 255. A brute force solution can divide an original pixel with grey level g into an 8×8 array of g black and 256-g white subpixels, and then encrypt each black and white subpixel separately by dividing it further into an array of subsubpixels with our techniques. However, we propose a more direct and elegant solution to the continuous tone visual encryption problem by using the following observation:

Each pixel in each one of the two transparencies is represented by a rotated half circle. When the two half circles (with rotation angles a and b) are carefully aligned, the superposition of the two half circles can range in colour from medium grey (representing white) to completely black (representing black) depending on the relative angle $a - b$ between the two rotated half circles (see Figure 3). If we choose for each pixel in each share a random absolute rotation angle (with the desired relative rotation angle between them), then each transparency will look uniformly grey and will reveal absolutely no information, but the superposition of the two transparencies will be a darker version of the original continuous tone image.

Another interesting extension of the original model deals with the problem of concealing the very existence of the secret message. Is it possible to send (by mail or fax) an innocent looking image of a house, superimpose on it an innocent looking transparency of a dog, and get a spy message with no trace of either the house or the dog? To construct such a scheme, we consider 2×2 arrays of subpixels, and define two types of shares (white with 2 black subpixels and black with 3 black subpixels) and two types of superimposed results (white with 3 black subpixels and black with 4 black subpixels). If the desired result is white, we use the shares presented in the top row of Figure 4 (along with their permutations). If the

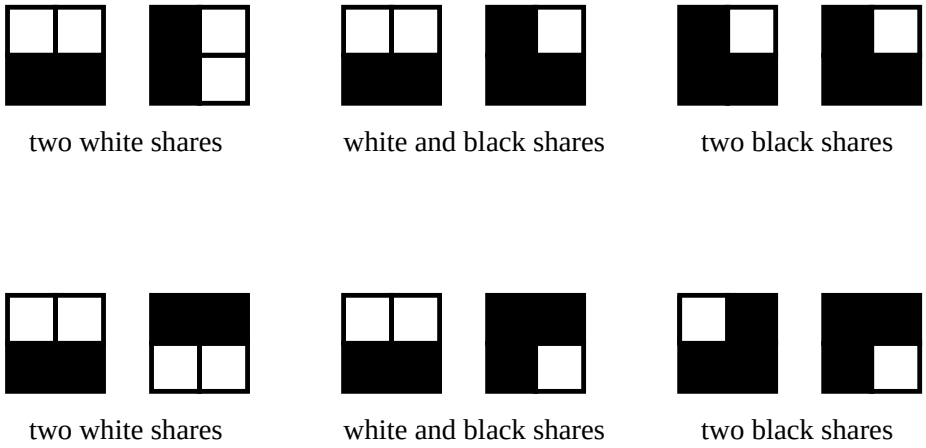

Figure 4: Use top row for white and bottom row for black

desired result is black, we use the shares presented in the bottom row of Figure 4 (along with their permutations):

The reader can easily convince himself that each transparency can contain an arbitrary image which reveals no information whatsoever about the superimposed image.

Acknowledgements

We thank Nati Linial for explaining his work on inclusion-exclusion, Ronny Roth and Mark Tuttle for careful reading of the paper and Ronen Basri for helping us with the figures.

References

[1] N. Alon, J. Bruck, J. Naor, M. Naor and R. Roth, *Construction of asymptotically good, low-rate error-correcting codes through pseudo-random graphs*, IEEE Transactions on Information Theory, 38 (1992), 509-516.

[2] N. Alon, O. Goldreich, J. Hastad and R. Peralta, *Simple constructions of almost k-wise independent random variables*, Random Structures and Algorithms **3** (1992), 289-304.

[3] N. Alon and J. Spencer, **The probabilistic method**, Wiley, 1992.

[4] J. L. Carter and M. N. Wegman, *Universal classes of hash functions*, Journal of Computer and System Sciences 18 (1979), pp. 143-154.

[5] J. Kahn, N. Linial and A. Samorodnitsky, *Inclusion-exlusion: exact and approximate*, manuscript.

[6] N. Linial and N. Nisan, *Approximate inclusion-exlusion*, Combinatorica **10**, 1990, pp. 349–365.

[7] F. J. MacWilliams and N. J. A. Sloane, **The theory of error correcting codes**, North Holland, Amsterdam, 1977.

[8] J. Naor and M. Naor, *Small bias probability spaces: efficient constructions and applications*, *SIAM J. on Computing*, vol 22, 1993, pp. 838–856.

[9] M. N. Wegman and J. L. Carter, *New hash functions and their use in authentication and set equality*, Journal of Computer and System Sciences 22, pp. 265-279 (1981).