

On the Security of a Visual Cryptography Scheme for Color Images^{*}

Bert W. Leung, Felix Y. Ng, and Duncan S. Wong

Department of Computer Science, City University of Hong Kong, Hong Kong, China
{bertleung,felix.citycs}@gmail.com, duncan@cityu.edu.hk

Abstract. In *Pattern Recognition*, vol. 36, 2003, Hou proposed a four-share visual cryptography scheme for color images. The scheme splits a secret image into four shares, the *black mask* and the other three shares. It was claimed that without knowing the *black mask*, no information about the secret image can be obtained even if all the other three shares are known. In this paper, we show that this may be true for a few specific two-color secret images only. In all other cases however, security cannot be guaranteed. We show that an attacker can compromise a randomly chosen two-color secret image from any two of the other three shares with probability $4/7$. The advantage will increase to $6/7$ if all the other three shares are known. If the secret image has three or four colors, we show that the attacker can compromise it with probability $4/7$ and $8/35$, respectively. Finally, we show that our technique can be extended to compromising secret images with more than four colors.

Keywords: visual cryptography, secret sharing, dithering, color decomposition, cryptanalysis

1 Introduction

Visual Cryptography Scheme (VCS), introduced by Naor and Shamir [1] in 1994, is a type of secret sharing [2] techniques for images. The idea of VCS is to split an image into a collection of random shares (printed on transparencies) which separately reveal no information about the original secret image other than the size of it. The image is composed of black and white pixels, and can be recovered by superimposing a threshold number of shares without any computation involved. Here is an example using a dithered black-and-white Lena image as the original secret image (Fig. 1).



Fig. 1. Original Secret Image - Dithered

By applying the Naor-Shamir 2-out-of-2 visual cryptography algorithm [1], two shares (printed on transparencies) are created, which separately reveal no information about the original image. It can

^{*} The work was supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China (RGC Ref. No. CityU 122107).

only be recovered when both of the shares are obtained and superimposed. Fig. 2 shows the two shares and the superimposition of them. Note that the size of the images is expanded by a factor of 4.




Fig. 2. Two Shares and the Superimposition of the Shares

The technology makes use of the human vision system to perform the OR logical operation on the superimposed pixels of the shares. When the pixels are small enough and packed in high density, the human vision system will average out the colors of surrounding pixels and produce a smoothed mental image in a human's mind. For example, a block of 2×2 pixels shown below will be viewed as a gray-like dot as the two black pixels and the two nearby white pixels are averaged out. If we print the 2×2 pixel blocks shown in Fig. 3 separately onto two transparencies and superimpose them. This effect is equivalent to performing a pixel-wise OR logical operation on each of the four pairs of pixels between these two transparencies. The result is shown in Fig. 4. One of the unique and desirable properties of VCS is that the secret recovery process can easily be carried out by superimposing a number of shares (i.e. transparencies) without requiring any computation.




Fig. 3. Two 2×2 pixel blocks



Fig. 4. Superimposed Image

Besides black-and-white images, a natural extension of this research problem is to perform secret sharing on color images. In [3], Hou proposed three VCS for color images. Among them, the first one uses four shares to split a secret image. The four shares are called *black mask*, *C (Cyan) share*, *M (Magenta) share* and *Y (Yellow) share*. This scheme reproduces the best quality among the three in terms of image contrast during secret image recovery process. It is also the only one supporting a

practically useful feature called *two-level security control*. This feature allows an authority to keep a particular share, the *black mask*, secret and release the other three shares to the public, without worrying about exposing the concealed image. In particular, the author claimed that this scheme is secure as long as the *black mask* is kept secret. There would have no information leaked even if all the other three shares, namely C , M , Y shares, are exposed regardless of the color composition of the original secret image.

Our Results. We conduct a security analysis on Hou’s VCS scheme with respect to the two-level security control feature. We find that the security of the scheme depends crucially on the color composition of the original secret image. We show that this scheme may support two-level security control only if the original secret image contains only two colors chosen from a specific set of colors we identified in this paper. If the original secret image contains any other colors, we find that an adversary will have a high chance of compromising the scheme. We propose an attacking technique and show that by applying this technique to Hou’s VCS scheme, an adversary will be able to recover the original secret image with high probability in the following scenarios.

- **Two-Color Case 1.** If the original secret image has two colors and any two of the C , M , Y shares are obtained, the adversary will be able to recover the original image with probability $4/7$.
- **Two-Color Case 2.** If the original secret image has two colors and all three of the C , M , Y shares are obtained, the adversary will be able to recover the original image with probability $6/7$.
- **Three-Color Case.** If the original secret image has three colors and all three of the C , M , Y shares are obtained, the adversary will be able to recover the original image with probability $4/7$.
- **Four-Color Case.** If the original secret image has four colors and all three of the C , M , Y shares are obtained, the adversary will be able to recover the original image with probability $8/35$.

In all the scenarios above, the probabilities are taken under the random choices of the colors in the original secret images. By recovering a secret image, we require that the adversary should at least be able to determine the shape or pattern of the original secret image, that is, being able to determine the boundary between two distinct color regions in the image.

We also show that our attacking technique can be extended to compromise images with more than four colors¹. It is noted that in any of the attacks, nothing about the *black mask* is needed to know by the adversary. This implies that the two-level security control cannot be guaranteed.

Paper Organization. In Sec. 2, we review Hou’s VCS. In Sec. 3, we formalize the scheme using matrix representation, then describe our attacks for the four scenarios above, and lastly discuss how to extend the attacking technique to scenarios where the original secret image has more than four colors. We conclude the paper in Sec. 4.

2 Review of Hou’s Four-Share VCS

2.1 Preliminary – Dithering

Dithering is a technique used to create an illusion of color depth in images with limited color palette. The technique has been used commonly on printing applications. In a dithered image, colors not available in the palette are approximated by a diffusion of colored pixels from within the available palette. The human eye perceives the diffusion as a mixture of colors within it. The basic principle of the diffusion is to pack pixels in higher density for representing darker colors and distribute the pixels sparsely for representing lighter colors.

¹ The original secret image can have at most 8 possible colors. This is because of the application of dithering on each of the three primitive color components. More details will be given in Sec. 2.2

Fig. 5 shows a continuous 256-level grey scale ramp, in which there are 256 levels (i.e. 256 distinct colors) of grey pixels with fixed distance apart from each other. Fig. 6 shows a 2-level dithered ramp after applying the Floyd-Steinberg dithering algorithm [4]. In the dithered image, the black dots are sparser in brighter parts and are denser in darker parts.



Fig. 5. A 256-level Grey Scale Ramp




Fig. 6. Dithered Ramp by Floyd-Steinberg Dithering Algorithm

2.2 Hou's VCS Scheme

In [3], Hou proposed three VCS for color images. Among them, there is one which produces four shares, namely *black mask*, *C* share, *M* share and *Y* share. Comparing with the other two schemes, this scheme yields the best image contrast during secret recovery which involves the superimposition of the four shares. In addition, this scheme is the only one which is claimed to support a practically useful feature called *two-level security control*. In the following, we review this scheme and illustrate it using an example where the original secret image is a 24-bit *color* Lena image shown in Fig. 7.



Fig. 7. Original Secret Image - 24-bit Color

Step 1. The scheme first decomposes the original image into three primitive-color images under the subtractive model, namely, *C* (Cyan), *M* (Magenta) and *Y* (Yellow). Fig. 8 shows the three primitive color components of the Lena image, where each image has 256 levels of the corresponding primitive color.

Step 2. After decomposition, each primitive-color image is dithered (e.g. by applying the Floyd-Steinberg algorithm [4]) so that each image will have two color levels, namely the presence of the




Fig. 8. Primitive Color (C , M , Y) Components - 256-Level

corresponding primitive color or the absence of it. Fig. 9 shows the three dithered primitive-color images and an illusion of their superimposed image. Each pixel in a dithered primitive-color image is having 1-bit color depth. The superimposed image is therefore having 3-bit color depth, that is, 8 colors altogether².



Fig. 9. Dithered C , M and Y Components and their Superimposition

Step 3. A *black mask* with double size of width and height of the original secret image is randomly generated. It contains the same number of *pixel blocks* as the number of pixels in the original secret image. Each block consists of four pixels which are aligned as a 2×2 square box with two black pixels and two ‘transparent’ pixels (when printed on a transparency). Since the *black mask* is randomly generated, for each pixel block, there are six possible patterns where Fig. 3 shows two of them.

Step 4. In this step, three other shares, namely C , M and Y shares are generated. To generate these shares, the dithered C , M and Y primitive-color images of the original secret image (e.g. Fig. 9) are read in pixel by pixel. As mentioned (in Step 2 above), each pixel in a dithered primitive-color image is now having only 1-bit color depth. If the value of the input pixel from a dithered primitive-color image is 1, that is, the pixel contains the primitive color, a pixel block in the corresponding primitive-color share will be produced. The pixel block produced will have the same pattern as the corresponding pixel block in the *black mask*.

For example, if the pixel at the upper left hand corner of the dithered C image of the original secret image is 1 and if the upper left hand 2×2 pixel block of the *black mask* is shown in Fig. 10, then the output pixel block shown at the upper left hand corner of the C share will have the same pattern as that in the *black mask*, as shown in Fig. 11.

On the other side, if the value of the pixel in the original dithered C image is 0, then the output pixel block of the C share will have the complement pattern of that in the *black mask*, that is shown in Fig. 12.

² In practice, one may start with a dithered image described above as the original secret image and carry out this VCS directly from Step 3.



Fig. 10. A *Black Mask* Pixel Block Example



Fig. 11. The *C* Share Pixel Block When the Original Dithered *C* Image is 1



Fig. 12. The *C* Share Pixel Block When the Original Dithered *C* Image is 0

Fig. 13 summaries this encoding method for an instance where a random pixel block of the *black mask* is chosen and shown in the first column of the Figure. The second column of the Figure shows the eight possible combinations of the original (dithered) *C*, *M*, *Y* pixel values. The following three columns show the encoding of the pixel blocks in the respective *C*, *M*, *Y* shares. The last column illustrates the superimposed image of *C*, *M*, *Y* shares with the *black mask*.

Mask	Original Pixel (C, M, Y)	Share 1 (C)	Share 2 (M)	Share 3 (Y)	Superimposed Pixel
	(0, 0, 0)				
	(1, 0, 0)				
	(0, 1, 0)				
	(0, 0, 1)				
	(1, 1, 0)				
	(0, 1, 1)				
	(1, 0, 1)				
	(1, 1, 1)				

Fig. 13. The Generation of *C*, *M*, *Y* Shares

Fig. 14 shows the four shares of the original Lena image and Fig. 15 illustrates the superimposition of these four shares. Similar to the Naor-Shamir visual cryptographic scheme for black-and-white images [1], the image size is expanded by a factor of 4.




Fig. 14. The *Black Mask* and the *C, M, Y* Shares

2.3 Two-Level Security Control

In this VCS, the purpose of the *black mask* is to cover up the ‘disguising’ but unwanted colors during the secret recovery process. Without possessing the *black mask*, it was believed in [3] that the ‘disguising’ colors could jumble with the actual colors of the original secret image and therefore could ensure the security of the original image. It was also claimed that without the *black mask*, the original image could remain secret even if all the three shares, i.e. *C, M, Y* shares, are known to an adversary.

In [3], this property is termed two-level security control and a practical application is given. For example, as long as a manager of a company keeps the *black mask* and gives the *C, M, Y* shares to his subordinates, the content of the image will remain confidential, even though all his subordinates collude. This could refer to the first security level control. The second level is that once the *black mask* is known, useful information about the original secret image could be revealed even if not all the *C, M, Y* shares are known. For example, Fig. 16 shows that useful information about the original Lena image is readily revealed by superimposing only the *black mask* and the *C* share.

3 Security Analysis

In this section, we analyze the security of the VCS reviewed in Sec. 2. We start off by formalizing the scheme using a matrix representation. As explained in Step 2 of the scheme (Sec. 2.2), after dithering,




Fig. 15. The Superimposition of the Four Shares



Fig. 16. The Superimposition of the *Black Mask* and *C* Share

there will be 8 distinct colors left in the original secret image. The following notations are used in our analysis.

- (C, M, Y) , where $C, M, Y \in \{0, 1\}$, represents the color of a pixel in the original dithered secret image in terms of cyan, magenta and yellow, respectively.
- $M_{(C, M, Y)}$ is a 4×4 Boolean matrix representing four pixel blocks which correspond to the same position in the *black mask* and *C*, *M*, *Y* shares, when the corresponding pixel color in the original dithered secret image is (C, M, Y) . The first row of $M_{(C, M, Y)}$ represents the pixel block of the *black mask*. The remaining three rows represent the respective pixel blocks of the *C*, *M*, *Y* shares. Each row consists of four values which are corresponding to the four pixels in the pixel block. These four values, from left to right, correspond to the upper left, upper right, lower left and lower right pixels in the pixel block. Among the four values, there must be two ones and two zeros. One indicates the presence of the corresponding color and zero indicates the absence of it. By extending the result of [1], we have the following fact.

Fact 1. If columns of a Boolean matrix are permuted, the patterns of the pixel blocks corresponding to the *black mask* and *C*, *M*, *Y* shares are changed accordingly. Note that

the illusion yielded after XORing all the four rows (which is equivalent to superimposing the *black mask* with C, M, Y shares) will remain the same for all the Boolean matrices generated by column permutation of $M_{(C,M,Y)}$.

- $\mathcal{C}_{(C,M,Y)}$ is the set of all possible Boolean matrices which correspond to a pixel in the original dithered secret image with color value (C, M, Y) . According to *Fact 1* above, we can see that suppose $M_{(C,M,Y)} \in \mathcal{C}_{(C,M,Y)}$, then a Boolean matrix generated from any column permutation of $M_{(C,M,Y)}$ is also in $\mathcal{C}_{(C,M,Y)}$. Also, for any $M_{(C,M,Y)} \in \mathcal{C}_{(C,M,Y)}$, the set of all distinct column permutations of $M_{(C,M,Y)}$ is $\mathcal{C}_{(C,M,Y)}$. This property is called **closed**. This can easily be shown by the method of contradiction and therefore is skipped here. Readers may refer to [1] for related details.

Fact 2. As each $M_{(C,M,Y)}$ consists of two pairs of identical columns, we can see that $\mathcal{C}_{(C,M,Y)}$ contains $\binom{4}{2} = 6$ Boolean matrices for each triple (C, M, Y) .

- $B_{(C,M,Y)}$ represents the *matrix basis* of the pixel color (C, M, Y) . It is a Boolean matrix chosen arbitrarily from $\mathcal{C}_{(C,M,Y)}$. Since $\mathcal{C}_{(C,M,Y)}$ is closed, given $B_{(C,M,Y)}$, $\mathcal{C}_{(C,M,Y)}$ is readily generated.

To share a secret pixel with color (C, M, Y) , the VCS (Sec. 2.2) randomly chooses a Boolean matrix $M_{(C,M,Y)}$ from $\mathcal{C}_{(C,M,Y)}$ and sets the pixel blocks of the *black mask* and C, M, Y shares to the values on the four rows of $M_{(C,M,Y)}$, respectively. In the following, we designate a matrix basis $B_{(C,M,Y)}$ for each pixel color (C, M, Y) as a Boolean matrix whose row for the *black mask* is $[1\ 1\ 0\ 0]$. Below are the 8 matrix bases³.

$$\begin{array}{ll}
 \text{White } B_{(0,0,0)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix} & \text{Yellow } B_{(0,0,1)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix} \\
 \\
 \text{Magenta } B_{(0,1,0)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \end{bmatrix} & \text{Magenta + Yellow } B_{(0,1,1)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix} \\
 \\
 \text{Cyan } B_{(1,0,0)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix} & \text{Cyan + Yellow } B_{(1,0,1)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix} \\
 \\
 \text{Cyan + Magenta } B_{(1,1,0)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \end{bmatrix} & \text{Black } B_{(1,1,1)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix}
 \end{array}$$

Before introducing the attacks, we give the following definitions.

Identical Pattern: In a Boolean matrix, two rows are having *identical pattern* if the positions of 1's and 0's are exactly the same.

Complemental Pattern: In a Boolean matrix, two rows are having *complemental pattern* if the positions of 1's in one row are exactly that of 0's in another row and vice versa.

³ One may choose other Boolean matrices for conducting the analysis.

Note that in any of the Boolean matrices for the VCS (Sec. 2.2), any two rows are either having identical pattern or complementary pattern.

In general, a secret sharing scheme should be secure regardless of the structure of the original secret, provided that the secret falls in the secret space of the scheme. When applied to the VCS reviewed in Sec. 2.2, we require that the security of the original secret image should not rely on the color composition (i.e. the number of colors) or pattern of it. If an adversary is able to determine at least the shape of the original secret image, that is, being able to determine the boundary between any two distinct color regions, from C , M , Y shares without knowing the *black mask*, then we consider the scheme to be insecure. This definition of security is commonly adopted and practically important because in practice, one can readily determine useful information of the original secret image if the shape of the image is revealed. For example, if the original image contains a well-recognizable logo or letters of some language alphabet.

In the following four subsections, we will evaluate four attacking scenarios as mentioned in Sec. 1. The attacking scenarios explore the differences on the color composition of the original dithered secret image and the number of shares that the adversary has obtained among the C , M , Y shares. In any of the attacks, the *black mask* is not known to the adversary, and therefore, they satisfy the two-level security control assumption. By referring to the application of two-level security control suggested in [3], all the attacking scenarios are related to collusion among a manager's subordinates who possess the C , M , Y shares, but not the *black mask*.

3.1 Two-Color Case 1

Below are the assumptions for this case.

- The original dithered secret image has **two colors**.
- Adversary \mathcal{A} has obtained any **two** of the C , M , Y shares.

We will show with high probability that \mathcal{A} is able to recover the shape of the original secret image. For all the probability evaluations below, they are taken under the random choices of the colors of the original dithered secret images. Without loss of generality, suppose the two shares obtained by \mathcal{A} are C and M shares. The following attacking technique can also be applied to cases if \mathcal{A} has obtained a different pair of shares.

Attack 1. After obtaining C and M shares, \mathcal{A} superimposes the shares and groups the pixel blocks into two categories: one having *identical pattern* and the other one having *complemental pattern*. Then \mathcal{A} ‘recolors’ the superimposed image by filling the pixel blocks in the group of *identical pattern* with one color and the pixel blocks in the group of *complemental pattern* with another color. Finally, \mathcal{A} checks the result and sees if there is any useful information such as the shape of the original secret image revealed.

Here is an example. Fig. 17 shows the original secret image IMG_A which contains two colors, magenta and yellow. Suppose \mathcal{A} has obtained the C and M shares (Fig. 18). When the two shares




Fig. 17. IMG_A : Original Secret Image - Two Colors




Fig. 18. C and M Shares of IMG_A

are superimposed, we can see in Fig. 19 that the shape of IMG_A is revealed. From the call-out image (1200% zoomed) in Fig. 19, we can see an obvious boundary between the region of *identical pattern* and that of *complemental pattern*. The region of *identical pattern* is the circle in the middle. Since the color of the circle in the original secret image is yellow, the pixel blocks of C and M shares in this region align with each other and hence have the *identical pattern*. This becomes clear when we refer to the Yellow matrix basis $B_{(0,0,1)}$. When cyan and magenta are superimposed, the resulting color shown is blue (please check the call-out image in Fig. 19). The region of *complemental pattern* is the exterior of the circle. Since the color of the original secret image is magenta. According to the Magenta matrix basis $B_{(0,1,0)}$, the C row and the M row in this matrix are having the *complemental pattern*. Therefore, we can see (in the call-out image) both cyan and magenta in this region.




Fig. 19. Superimposed Image of C and M Shares of IMG_A

After applying the ‘recoloring’ described in the paragraph “*Attack 1*” above, we can completely recover the shape of the original secret image (Fig. 20). As there are only 8 possible colors in the original secret image, there are altogether only 56 ways of coloring the two regions and one of them is the original secret image.

Analysis. We now evaluate how likely that \mathcal{A} can compromise the scheme in this case. From the two shares obtained by \mathcal{A} , as explained above, there are only two possible outcomes when these two shares are superimposed. One outcome corresponds to the *identical pattern* of the pixel blocks and the other corresponds to the *complemental pattern* of the pixel blocks at the corresponding position




Fig. 20. Recolored Superimposition of C and M Shares of IMG_A

of the two shares. If the pixel blocks are having *identical pattern*, then among the 8 matrix bases, the following four matrix bases satisfy the condition that the rows corresponding to C and M shares are having *identical pattern*.

$$\begin{array}{ll}
 \text{White } \mathbf{B}_{(0,0,0)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix} & \text{Yellow } \mathbf{B}_{(0,0,1)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix} \\
 \text{Cyan + Magenta } \mathbf{B}_{(1,1,0)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \end{bmatrix} & \text{Black } \mathbf{B}_{(1,1,1)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix}
 \end{array}$$

Let G_1 be the set of these four matrix bases, namely, $G_1 = \{\mathbf{B}_{(0,0,0)}, \mathbf{B}_{(0,0,1)}, \mathbf{B}_{(1,1,0)}, \mathbf{B}_{(1,1,1)}\}$. If the pixel blocks in the corresponding C , M shares are having *complemental pattern*, then the corresponding pixel color of the original secret image could correspond to one of the following four matrix bases.

$$\begin{array}{ll}
 \text{Magenta } \mathbf{B}_{(0,1,0)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \end{bmatrix} & \text{Magenta + Yellow } \mathbf{B}_{(0,1,1)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix} \\
 \text{Cyan } \mathbf{B}_{(1,0,0)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix} & \text{Cyan + Yellow } \mathbf{B}_{(1,0,1)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}
 \end{array}$$

Let $G_2 = \{\mathbf{B}_{(0,1,0)}, \mathbf{B}_{(0,1,1)}, \mathbf{B}_{(1,0,0)}, \mathbf{B}_{(1,0,1)}\}$. As there are two colors in the original secret image. Suppose that one of these two colors is in G_1 (the group of *identical pattern*) and the other color is in G_2 (the group of *complemental pattern*). As we have seen in the example above, since the color display of the superimposed C and M shares for a color from G_1 is different from that for a color from G_2 , it is easy for adversary \mathcal{A} to find the boundary between two color regions. Let \mathbf{E} be the event that in a two-color original secret image, one color is from G_1 and the other from G_2 . As the cardinality

of each of G_1 and G_2 is 4, there are $4 \times 4 = 16$ possible combinations for event \mathbf{E} to happen. If we consider an experiment where the two colors in the original secret image are randomly chosen from the 8 possible colors, we can see that

$$\Pr[\mathbf{E}] = \frac{16}{\binom{8}{2}} = \frac{4}{7} \quad (1)$$

which is the success probability of \mathcal{A} in breaking the scheme in this attacking scenario.

3.2 Two-Color Case 2

Below are the assumptions for this case.

- The original dithered secret image has **two colors**.
- Adversary \mathcal{A} has obtained all three shares, namely the C , M and Y shares.

When comparing with the assumptions in Two-Color Case 1, \mathcal{A} in this case has more information and is more capable. This is because \mathcal{A} can simply launch “*Attack 1*” (page 10) on any two of the shares it gets and the success chance will at least be equal to that in the Two-Color Case 1, that is, $4/7$. In the following attack, we will see that \mathcal{A} can do a little bit more to boost its success probability in a significant way.

Attack 2. \mathcal{A} superimposes all three shares and categorizes the superimposed pixel blocks into two groups, where in each group, the pixel blocks are having the same *mixture of colors*, while between these two groups, they are having different mixture of colors. If such two groups are found, \mathcal{A} ‘recolors’ the superimposed image by filling the two groups separately with two distinct colors.

We use the same example as in Sec. 3.1 to illustrate this attack. The original secret image is IMG_A (Fig. 17). The C and M shares have been shown before in Fig. 18, and the Y share is shown in Fig. 21. After superimposing C , M , Y shares, we obtain the image shown in Fig. 22. From the call-out image (1200% zoomed) in Fig. 22, we can see the two regions clearly. The two groups of different mixture of colors can also be identified easily. After recoloring (Fig. 23), the shape of the original secret image is recovered.

Analysis. Similar to the analysis in Sec. 3.1, we consider the groups of matrix bases which give distinct colors when C , M and Y shares are superimposed. It can be found that there are four groups, each of the groups gives a distinct *mixture of colors* when C , M , Y shares are superimposed. Below are the groups.

Group 1 – {Black,White}. All three shares are having the *identical pattern*. The mixture of colors when C , M , Y shares are superimposed contains Black and White.

$$\text{White } \mathbf{B}_{(0,0,0)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix} \quad \text{Black } \mathbf{B}_{(1,1,1)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix}$$

Group 2 – {Blue,Yellow}. C and M shares are in the *identical pattern* but not Y share. The mixture of colors contains Cyan+Magenta=Blue and Yellow.

$$\text{Yellow } \mathbf{B}_{(0,0,1)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix} \quad \text{Cyan + Magenta } \mathbf{B}_{(1,1,0)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \end{bmatrix}$$




Fig. 21. Y Share of IMG_A




Fig. 22. The Superimposition of C, M, Y Shares of IMG_A

Group 3 – {Green,Magenta}. C and Y shares are in the *identical pattern* but not M share. The mixture of colors contains Cyan+Yellow=Green and Magenta.

$$\text{Magenta } \mathbf{B}_{(0,1,0)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \end{bmatrix} \quad \text{Cyan + Yellow } \mathbf{B}_{(1,0,1)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}$$

Group 4 – {Red,Cyan}. M and Y shares are in the *identical pattern* but not C share. The mixture of colors contains Magenta+Yellow=Red and Cyan.

$$\text{Cyan } \mathbf{B}_{(1,0,0)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix} \quad \text{Magenta + Yellow } \mathbf{B}_{(0,1,1)} : \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix}$$

Each group above has a distinct mixture of colors when all the C, M and Y shares are superimposed. We refer to these groups as **color mixture groups**.

If the two colors of the original secret image come from two different color mixture groups, then the shape of the original secret image will be revealed. If these two colors are chosen randomly from the 8 possible colors, the probability that \mathcal{A} breaks the VCS in this attacking scenario will be $1 - \frac{4}{\binom{8}{2}} = \frac{6}{7}$, where $\frac{4}{\binom{8}{2}}$ is the probability that the two colors belong to the same color mixture groups.




Fig. 23. Recolored Superimposition of C , M , Y Shares of IMG_A

3.3 Three-Color Case

We now investigate the case where the original secret image has three colors. Below is the summary of the assumptions for this case.

- The original dithered secret image has three colors.
- Adversary \mathcal{A} has obtained all three shares, namely the C , M and Y shares.

Since \mathcal{A} gets the same set of shares as in the Two-Color Case 2, if \mathcal{A} superimposes all the three shares and categorizes the pixel blocks with respect to the four color mixture groups (Sec. 3.2), \mathcal{A} will be able to get *up to* three groups of color mixtures. By recoloring these identified color mixture groups with distinct colors, the shape of the original secret image could be recovered.

Suppose the original secret image is IMG_B (Fig. 24) which contains three color stripes. Fig. 25 shows the C , M , Y shares of IMG_B . After superimposing the shares, the image obtained is shown in Fig. 26. We can see clearly from the call-out images (1200% zoomed) the boundary between any two adjacent color regions. Fig. 27 shows the superimposed image after recoloring.




Fig. 24. IMG_B : Original Secret Image – Three Colors

Analysis. In order to recover all the three colors of the original secret image, these three colors should fall into three distinct color mixture groups. There are $\binom{4}{3} = 4$ ways of picking three distinct color mixture groups. As there are two colors in each color mixture group, there are altogether $4 \times 2^3 = 32$ number of possible combinations of the three colors in the original secret image for the attack to work. Therefore, the success probability of \mathcal{A} is $\frac{32}{\binom{8}{3}} = \frac{4}{7}$.

3.4 Four-Color Case

We now extend our attacking technique to the case where the original secret image has four colors. Below is the summary of the assumptions for this case.




Fig. 25. C, M, Y Shares of IMG_B




Fig. 26. Superimposed Image of C, M, Y Shares of IMG_B

- The original dithered secret image has four colors.
- Adversary \mathcal{A} has obtained all three shares, namely the C, M and Y shares.

The attack is the same as below except that in this case, \mathcal{A} may get *up to* four color mixture groups from the superimposed pixel blocks. Suppose the original secret image is IMG_C (Fig. 28) which contains four colors. Fig. 29 shows the C, M, Y shares of IMG_C . After superimposing the shares, the image obtained is shown in Fig. 30. We can see from the call-out images (1200% zoomed) the boundary between any two adjacent color regions. Fig. 31 shows the superimposed image after recoloring.

Analysis. Similar to the analysis in Sec. 3.3, \mathcal{A} can recover the original secret image if its four colors fall into four distinct color mixture groups given in Sec. 3.2. Similarly, we can find that the success probability of \mathcal{A} in this case is $\frac{\binom{4}{2} \times 2^4}{\binom{8}{4}} = \frac{8}{35}$.

3.5 Remarks

From the attacks above, readers may find that \mathcal{A} will not be able to compromise the secret if the original secret image contains only one of the following four color pairs:

- {White, Black}
- {Cyan + Magenta, Yellow}
- {Cyan + Yellow, Magenta}
- {Magenta + Yellow, Cyan}




Fig. 27. Recolored Superimposition of C , M , Y Shares of IMG_B




Fig. 28. IMG_C : Original Secret Image – Four Colors

For each of the color pairs above, if we consider their matrix bases but ignoring the first row, which is corresponding to the *black mask*, their matrix bases are identical to each other after conducting some column permutation to one of them. This implies that no matter how the C , M , Y shares are superimposed, the superimposed pixel block will consist of the same mixture of colors. Therefore, the VCS reviewed in Sec. 2.2 may be secure if only these color pairs are used. This however, severely limits the effectiveness of the scheme as it may not be used to perform secret sharing for secret images which have three colors or more without worrying about leaking some useful information to colluders who possess the C , M and Y shares.

3.6 Attack Extension

From the color mixture groups given in Sec. 3.2, we know that if the original secret image has more than four colors, then there must exist at least two colors in the original secret image that belong to the same color mixture group. However, if the regions of these two colors are not adjacent to each other, attackers may still be able to recover the shape of the original secret image by superimposing the C , M and Y shares. We illustrate this by showing an example below.

Fig. 32 shows a 5-color secret image denoted by IMG_D and Fig. 33 shows the C , M and Y shares of IMG_D . After superimposition, five color regions could be identified (Fig. 34) and they become clear after the recoloring is performed (Fig. 35). Although the adversary \mathcal{A} may not be able to tell if the two diagonal regions painted with cyan in Fig. 35 correspond to two different colors in the original secret image (as \mathcal{A} may not know the number of colors in the original image), the shape of the original secret image is completely recovered.

However, if the regions of two colors which belong to the same color mixture group are adjacent to each other, \mathcal{A} is not able to find out the boundary between these regions. We illustrate this using the following example.

Fig. 37 shows the C , M and Y shares of some secret image. After superimposing them, there are four distinct color regions identified (Fig. 38). It becomes clear in the recolored image shown in




Fig. 29. C, M, Y Shares of IMG_C




Fig. 30. Superimposed Image of C, M, Y Shares of IMG_C

Fig. 39, which is not exactly the same in shape or pattern as the original secret image which is shown in Fig. 36.

4 Conclusion

In this paper, we provided a detailed security analysis to Hou's four-share VCS [3] which was previously believed to support the two-level security control. We found that by superimposing the shares that an adversary has acquired and applying some recoloring method according to the color mixture groups we identified (Sec. 3.2), the adversary may find out some useful information such as the shape or pattern of the original secret image. We provided some examples to support our findings and evaluated the adversary's success probability under various settings. In particular, the adversary can find out the shape of the original secret image with probability $4/7$ if the original secret image has two colors only and the adversary has acquired two of the C, M and Y shares. The winning chance of the adversary increases to $6/7$ if the adversary gets all the three shares. If the original secret image has three or four colors, we showed that the adversary will still be able to compromise the secret with winning probability $4/7$ and $8/35$, respectively. Nothing about the *black mask* is needed to know by the adversary in any of the attacks. We further discussed the feasibility of our attack against secret images which have more than four colors. We illustrated the case where the adversary will still be able to recover the shape of the secret image entirely and the other case where the adversary may only be able to retrieve partial information of the original secret image.

Our results suggest that the security of the scheme depends critically on the color composition and distribution of the original secret image. If the original secret image is composed of two specific colors




Fig. 31. Recolored Superimposition of C , M , Y Shares of IMG_C




Fig. 32. IMG_D : Original Secret Image – Five Colors

given in Sec. 3.5, the scheme may be secure. Otherwise, the scheme is highly susceptible to attacks which can effectively compromise the scheme in various scenarios.

References

1. M. Naor and A. Shamir. Visual cryptography. In *Advances in Cryptology - EUROCRYPT '94*, pages 1–12, 1994. Lecture Notes in Computer Science, Vol. 950. (Cited on pages 1, 6, 8, and 9.)
2. A. Shamir. How to share a secret. *Communications of the ACM*, 22:612–613, November 1979. (Cited on page 1.)
3. Y. C. Hou. Visual cryptography for color images. *Pattern Recognition*, 36:1619–1629, 2003. (Cited on pages 2, 4, 7, 10, and 18.)
4. R. W. Floyd and L. Steinberg. An adaptive algorithm for spatial grey scale. In *Proc. the Society of Information Display*, volume 17, pages 75–77, 1976. (Cited on page 4.)




Fig. 33. C, M, Y Shares of IMG_D




Fig. 34. Superimposed Image of C, M, Y Shares of IMG_D




Fig. 35. Recolored Superimposition of C, M, Y Shares of IMG_D



Fig. 36. IMG_E : Original Secret Image – Five Colors




Fig. 37. C, M, Y Shares of IMG_E

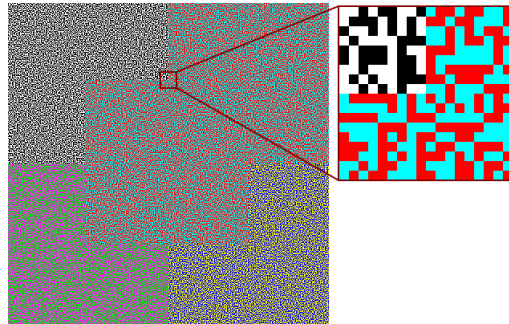


Fig. 38. Superimposed Image of C, M, Y Shares of IMG_E




Fig. 39. Recolored Superimposition of C, M, Y Shares of IMG_E