Chapter 7

Suspicion and other issues concerning the secure image database model

This chapter continues the discussion regarding secure image databases by identifying some requirements and issues that may need to be adhered to in order to prevent potential problems that may occur for certain implementations of the model. Suspicion is discussed more extensively than in the previous chapter, and possible solutions are provided which should help reduce suspicion. Some extensions to the model will then be presented which should produce more accurate region selections. Finally, important security issues concerning the model will be discussed.

7.1 Suspicion

The secure image database models presented in the previous chapter must at some point assemble an image according to a security classification, either in the first phase (by the database interface) or in the third phase (by the SMI). Unauthorised parts of the original image must be removed during this process so that only authorised parts of the image remain (for that particular security classification). If unauthorised parts of an image are removed and only the remainder of the original image is simply transmitted to the entity that requested the image, the requesting entity will probably *suspect* that the original image has been tampered with. Figure 7.1 demonstrates this potential problem.

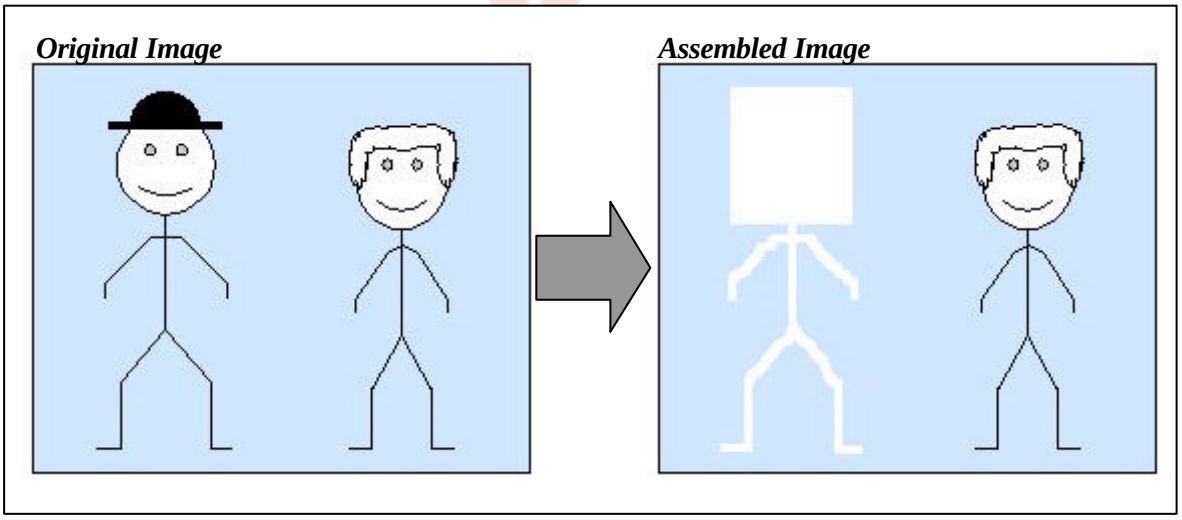

Figure 7.1: Basic illustration of suspicion

Figure 7.1 shows a simple example of an original image (on the left) containing a stick-man and a stick-woman on a light blue background. The man is selected using the database interface and is associated with an image object; let us call this image object ‘S-man’. The image on the right displays the assembled image after a request has been made by an entity whose security classification is below the classification assigned to image object ‘S-man’. Looking at the assembled

image, it is obvious that a requesting entity will suspect that the original image has been modified (and in this case also where the changes have been made).

The goal when dealing with the issue of suspicion is therefore to minimise the probability that the receiving entity will come to the conclusion that changes were made to the original image. Exactly how important it is to ensure this depends very much on the sensitivity of the multimedia data, the multimedia application making use of the data, and the organisation(s) concerned. It may not however be very important to prevent complete suspicion (i.e. to ensure that the receiving entity has no idea of any modifications that have been made); it may only be necessary to prevent the entity from knowing *where* in the assembled Image certain alterations have been made. Further, different degrees of importance may be required for each security classification.

7.1.1 Factors that raise the level of suspicion

The level of detail in an image varies from simple line drawings to complex impressions of photographic quality. It is generally harder to conceal the changes made to an image containing detailed worldly objects such as people, animals, food, and flowers. The background for an image also plays an important part in determining the probability of suspicion. It is more likely to suspect an image containing a detailed background that has many curves and colours, simply because the removal of worldly objects within such images leave notable 'holes' or empty regions within the assembled image.

There are a number of factors that increase the level of suspicion for a particular image. Each modification that needs to be made to an image generally heightens the level of suspicion. Overlapping regions can cause problems, especially if only one of the regions is associated with an unauthorised image object. It can be very difficult to remove the one region without removing a small part of the other, especially for detailed regions. On the other hand, a small piece of the region associated with the unauthorised image object may be left attached to the overlapping region.

It is also very difficult to correctly select a worldly object if its surface reflectivity has fine-grained variation (called *visual texture*), such as a patch of grass. This is mainly due to the fact that a large number of small regions of substantially different intensity make up the worldly object. The automated method cannot therefore be used to select this region because the variation in intensity is too

large. The patch of grass may for example surround other detailed worldly objects that must not be selected, preventing us from using the box or polygon method to select the patch of grass.

Another factor which can increase the level of suspicion is the occurrence of random pixels of contrasting intensity within a region, which may happen when converting a high quality image to a lower quality (averaging may not always remove this). If the region is replaced with a constant background and these pixels are not also selected with the region, this could leave obvious signs in the assembled image that it has been changed, and may even allow the requesting entity to deduce what areas of the image have been altered. It is also preferable (when aiming to reduce suspicion) to be able to select a lesser amount of regions associated with a worldly object. This is because it is less likely that an entity would 'miss' associating a region that *actually* belongs to the object.

The impact of these factors varies on a case-by-case basis; the more often they occur, the higher the probability that an entity will suspect changes to have taken place. The amount of effort needed to conceal all changes made to an image also varies from one image to another. Most of these problems occur because accurately selecting the regions of a worldly object can be very difficult at times. The entity defining the objects must not leave 'clues' for entities that are not authorised to view an image object, such as pieces of a worldly object whose regions should *all* be associated with that image object. Thus improving the methods used to select these regions should help reduce suspicion in general.

7.1.2 Lowering the level of suspicion with background images

The solution provided in the models that were presented in the previous chapter was to make use of a 'background image'. This image ideally consists of the basic scenery of the original image, and does not contain the worldly objects that would most likely be assigned to an image object. Any region removed from the original image during the assembly process can then be replaced with the corresponding area of the background image. A landscape setting like the ocean or a beach (without the people) is an example of an ideal background image for an original image that has this as its background setting.

The SMI or database interface (if the images are pre-calculated) copies the relevant parts of the background image to 'cover-up' the suspicious parts in the resulting image. If the actual background for the original image is not available,

another background image that fits reasonably in place of the original scenery should be used. As mentioned earlier, if no such image is available, a default, uniform coloured background image may be used instead. Figure 7.1 above should for example be associated with a light-blue background image of uniform intensity. The unauthorised regions of the image will then be replaced with the light blue colour of the background, which should reduce suspicion.

Certain cases do arise where using an ideal background image still does not completely solve the suspicion issue. This happens when two image objects each associated with a different worldly object have different security classifications, the image object with higher classification is overlapping the other image object, and only the regions covered by the image object with higher classification must be replaced (using a part of the background image). This is because the overlapping area is now part of the scenery, whereas it should actually be the missing part of the authorised worldly object. Figure 7.2 and 7.3 illustrate this problem.

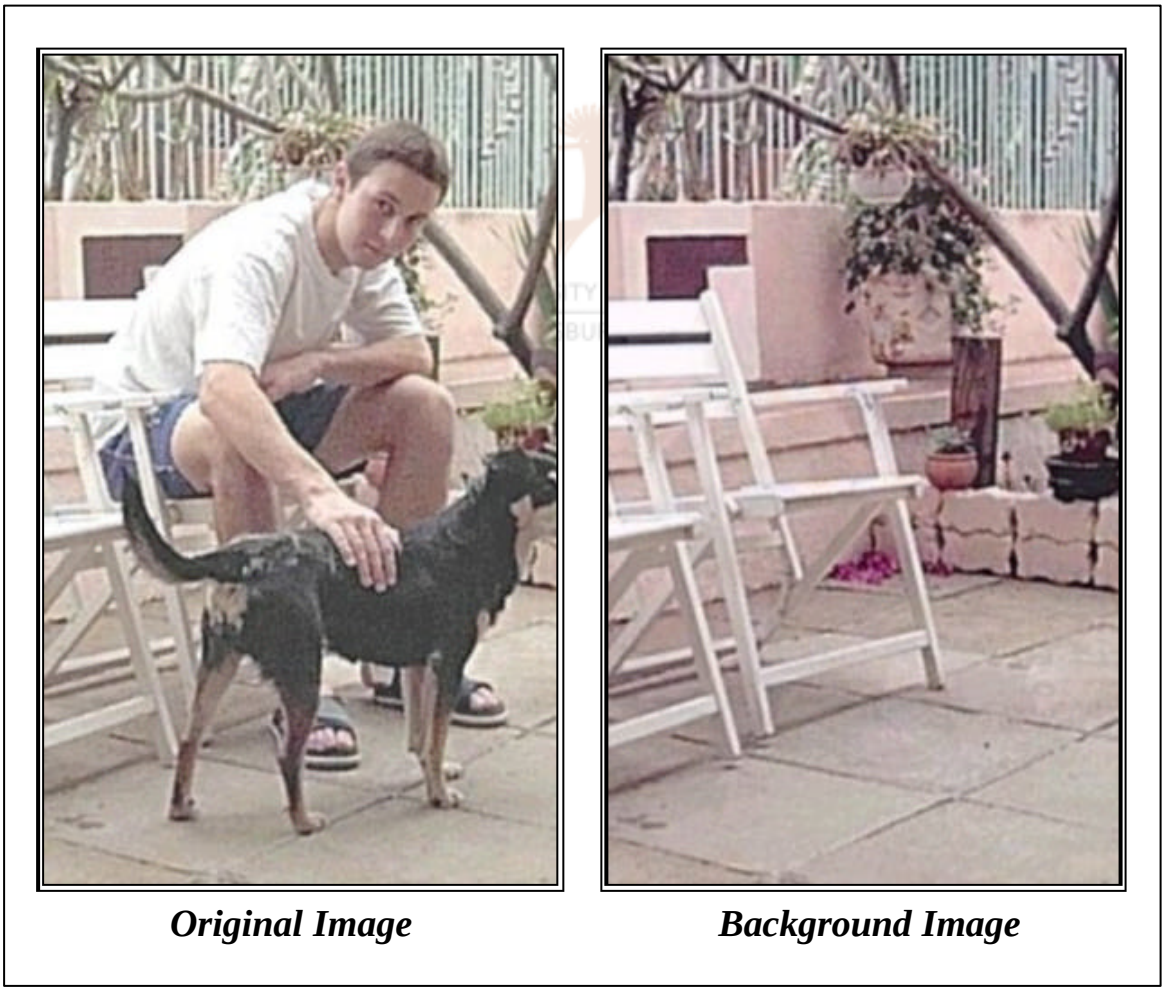

Figure 7.2: Overlapping image object problem (original and background images)

Figure 7.3: Overlapping image object problem (assembled image)

Figure 7.2 consists of two photographs. The photo on the left is the original image of a boy (Peter) and his dog (Rocky) standing in front of him. The photo on the right displays the background of the original image which excludes the worldly objects Peter and Rocky. The secure image database model is used to classify the areas of this original image. In this example, Rocky was selected using the polygon and automated methods and assigned the highest possible classification, 'General'. The worldly object Peter was not classified, and so his associated classification is 'Private' (the lowest classification). The image in figure 7.3 was generated using the secure image database model prototype after a request was made using the security classification 'Private'. Notice that a part of Peter's legs and feet are now missing due to the overlapping image object problem.

7.1.3 Other methods of lowering the level of suspicion

The level of suspicion may be dramatically reduced by adding objects to the assembled image when requested by an entity with a given security classification.

The object to be added may be a worldly object that helps to give more meaning to the resulting image after the unauthorised parts have been removed (the assembled image therefore makes more sense with these additions). A typical example of this is an image of a person leaning on another person who for the requesting entity's security classification has been regarded as unauthorised. The person being leant on will then be removed and replaced with the background scenery, which may not unfortunately make gravitational sense to have the remaining person leaning on nothing.

This can be solved by adding another person or object in place of the 'unauthorised person' whenever he/she must be removed from the image. It may often be the case that only a part of a worldly object needs to be added to make an image more believable. For example, only a part of Peter's legs and feet in figure 7.3 must be added to "complete the picture". Note that a worldly object or a part of one must be added after the background image is used to replace all the unauthorised areas (or else the additional object may be partially or even entirely substituted by the corresponding part of the background image).

It may sometimes be the case that a number of possible security classifications have some common image objects shared between them, although different regions must be used to replace the unauthorised areas for each of these classifications. An example of this is if an original image displaying a section of a military camp must be assembled in a way that the requesting entity believes that it is his/her barracks that forms part of the image i.e. depending on the rank or classification level of the requesting entity. This problem can be solved by adding the relevant barracks as a worldly object to the selected position within the image, although this could be difficult to manage if many worldly objects need to be added/replaced depending on the security classification.

A better solution may therefore be to associate a different background image with the image identifier for each available security classification. This way, any part of the background can be required to change in order to suit the needs of each classification without adding a list of worldly objects for each case. This method can therefore easily be used to completely change how the assembled image will look like if required. The secure image database model will need to accommodate this by allowing the entity to choose a list of background images (instead of just one) and associate each of these with a classification. This will form part of the image identifier by simply adding these new fields/attributes and recording the extra background images in the repository. The assembly process will also need to be slightly modified to use the relevant background image.

A white or uniform coloured background that is used as a background image may not always be suitable, especially when preventing suspicion is a big concern. Another method that can be tried in these cases is to replace the unauthorised area with the surrounding pixels instead of with the corresponding area of a background image, which should be especially effective for small regions. The intensity of an adjacent pixel or an average of the nearby *x* number of pixels that are a specified upper bound distance away can be used to replace the missing pixels. The texture of the surrounding area could even be analysed using statistical techniques, and a similar texture could be used to fill up the missing area. An example of this is when a person is removed from a patch of grass, requiring us to replace the unauthorised area with similarly textured grass. See Appendix A for more examples regarding the issue of suspicion.

7.2 Region selection improvements

There are a number of reasons encouraging us to improve the region selection process. Most importantly, by improving selection we aim to:

- select more accurately
- make it easier to select the desired region
- select the regions more efficiently
- help reduce suspicion

It is ideal to minimise the number of regions associated with one specific image object (assuming that the image object fully represents a worldly object). This is advantageous to all three phases of the secure image database model. Firstly, less effort is required to select the desired worldly object (phase 1). Secondly, less data needs to be written to the image database (phase 2). Thirdly, less data needs to be retrieved, and a smaller number of regions associated with the image object need to be reselected when the image is assembled (phase 3 for the original model and phase 1 for the model variation).

Any improvement made to the selection process should reduce processing requirements in general for the server(s) running the database interface and the SMI components. These improvements should also help reduce suspicion. A more accurate method of selection will minimise the likelihood of not selecting parts of unauthorised regions, which can easily happen if the selection process is rushed (especially for very detailed images). Similarly, it will help prevent areas of authorised worldly objects from being included in the selection. Finally, an image

that takes a long time to assemble may also arouse suspicion. More efficient selection (and reselection) will help prevent this.

The easiest way of improving the selection process is to increase the number of region selection methods available to the entities that are creating the image objects. This allows the entities more choice in deciding which method is best suited to selecting the desired worldly object. An entity also has a wider range to experiment with in order to ensure that the selection requirements are satisfied, depending on which factors are the most important i.e. accuracy, efficiency, suspicion etc.

The algorithms used to calculate the regions should also be optimised and if possible improved to yield better results. The automated method discussed earlier could for example be modified to select *all* the pixels within the boundary that is calculated using the specified parameters i.e. a second type of automated method which calculates the outer border and selects all pixels within that border. This algorithm modification should help to select regions containing a certain amount of noise more accurately than the original algorithm. Small regions that must not be selected within larger regions would however also be selected using this method if the outer border was calculated for the larger region e.g. selecting a car's tyre would result in the selection of the entire wheel. Another variation could also compare pixels within a region to a uniform colour (as opposed to neighbouring colours). No matter what improvement is made, it must be ensured that all new information required to reselect the regions is available at the time of assembly.

Another way of improving region selection is to allow the methods to be combined when necessary. Many cases occurred while using the prototype to select a region with the automated method where even though the desired region was properly selected, parts of unwanted regions would be selected together with the desired region. This problem can be solved by allowing the automated method to be used together with another method; the polygon method can for example be used to approximate the outer limit of the final selection while the automated method is used to actually select the region up to the outer limit i.e. all the area selected using the automated method up to the boundary selected using the polygon method is counted as one region.

Similarly, two methods can be combined in such a way that the selected region consists of all the area selected using the one method and *excludes* all the area selected with the second method within the first selection (like a 'NOT operation'). The polygon method may therefore be used to select a large area while the

automated method is used to select small regions within that large area that must not be included in the final region selection. This can definitely save time in cases where most of the regions within an area must be selected (which can be selected using a single region selection operation), aside from a few regions within that area that can also easily be selected (separately). It is therefore easier in these cases to identify what must be *excluded* within certain boundaries than to select all the regions that must be *included*.

7.3 Security issues

Although the above secure image database models are aimed at ensuring logical access control for the content within specified images that are stored in an image repository, a number of security requirements must first be satisfied in order for a system based on these models to function securely. Some of these requirements may be more important than others; this depends very much on the actual physical implementation of the model and the sensitivity of the multimedia data e.g. it may not be necessary to ensure the integrity of the data transferred between the database interface and the RDBMS if these components are on the same server. Each phase within the models presented has its own requirements, which will now be briefly discussed.

7.3.1 Security requirements for phase 1

Firstly, we must ensure that all the image objects can only be created and managed using the database interface. This is related to the next security requirement: only the database interface is allowed to have direct access to the image database. A secret key can for example be exclusively shared between these two components, in which case it must be guaranteed that only correctly encrypted data is transferred from the database interface to the image database.

If the entity defining the objects is in a remote location (away from the server containing the secure image database), two implementation options exist. The first option is to place the database interface on the client computer. Once the entity has finished defining all of the image objects, the database interface must securely transfer all the region and image object data and the original image to the image database. The second option is to place the database interface on the server. In this case, the client computer must contain a secure program whose function is to remotely communicate with the database interface in a secure

manner. The entity sends commands to the database interface via the client program, while the server performs the necessary calculations and sends back the results to the client (a more powerful server is therefore needed in this case).

As mentioned in the previous chapter, the DAC or MAC policy defined in the security policy determines which entity can create image objects. Similarly, the DAC or MAC policy can be used to determine which entity/entities are authorised to manage the defined image objects. For example, it may be the case that only the entities that have a security classification equal to or above an image object's classification are allowed to change region specific information or classifications associated with that image object.

7.3.2 Security requirements for phase 2

The database interface must ensure that all data is securely transferred to the image database. The image database must in turn verify that it is really the database interface that is requesting the storage or retrieval of the image objects, and that all communication with the database interface is also encrypted. Similarly, the image database must verify that it is the SMI that is requesting the original image and its associated information, and that all communication with the SMI is also encrypted.

Exclusive communication must be ensured between the database and the image repository. These two components could be configured as an internal network that is not directly accessible through the public network. The server containing the database interface could then be configured as a proxy server through which incoming requests must be made. If necessary, the entire database could be encrypted and only the requested data decrypted at the time of the request (although this will delay the assembly process even further). Higher security applications may also require that the image repository encrypts all files before storage and decrypts requested image files before transferral to the RDBMS.

7.3.3 Security requirements for phase 3

The SMI must at all times ensure that the requesting entity cannot access the image database directly, since this will allow the entity to directly make a request to the RDBMS for the original image (even though the RDBMS should not accept requests that have not been encrypted with the shared key). If the SMI is on the

requesting entity's computer, it must be stable enough to ensure that the original image cannot be accessed during the assembly process, especially since this process may not be very efficient, giving the entity enough time to try and hack into the program and access the original image (e.g. while it is still in memory).

If possible, the SMI should be located on the server and not the client machine, since this is more secure (even though the system requirements will be made higher on the server). In this case, a client program will be used which will receive the encrypted image that has been assembled, decrypt it, and make it accessible to the requesting entity. This approach is more secure because the image is transferred to the entity only once it has been assembled at the server's side i.e. the client machine does not receive the original image in any form (preventing an entity from receiving the original image through a program hack). This approach also helps prevent suspicion because at times when assembling the image is not efficient, the entity may just think that it is the network that is lagging (obviously within reasonable time limits) – if the SMI is on the client machine, the requesting entity may determine when the encrypted original image has been received which will increase suspicion for a long assembly process.

7.3.4 Security requirements in general

As with most computer systems, the security of an application based on the secure image database model depends on proper identification and authentication of its entities. An entity creating the image objects for a particular image must be properly authenticated to have the claimed identity before checking whether the entity is authorised to create/modify image object information for that image. Similarly, a requesting entity must be properly authenticated before requesting an image that will be assembled according to its security classification.

Communication between the different components of the model will most of the time be encrypted to ensure the confidentiality of the data i.e. data sent between the entity and the database interface, the database interface and the database, the SMI and the database, and the SMI and the requesting entity (encrypted both ways for all of these). Some of these communication lines may however not need to be encrypted for certain implementations e.g. the server may be secure enough to avoid encrypting data between the database interface and the image database if they both reside on the server. The type and strength of the encryption used also depends on the implementation and may vary from one communication line to another. As already mentioned in chapter 2, symmetric key-based encryption

should mostly be used for more efficient encryption/decryption (public key encryption should only be used to exchange session keys).

Although not as essential as the confidentiality of the data (for this model), the integrity of the data sent between each model component also depends on the implementation and the components in question. More emphasis may for example be placed on the integrity of the data sent between the SMI and the requesting entity because it is most likely that an attacker would modify the data between these two components. The IS service of non-repudiation may be considered least important, except in cases where an entity must be held accountable for changing an image object's associations, or where a requesting entity may need to be held accountable for an image request. Lastly, the availability of the data should always be ensured e.g. the database interface and SMI must be reliable and stable enough to prevent DoS attacks whether on the server or client.

7.4 Conclusion

Reducing the estimated level of suspicion for an image works on a case-by-case basis. Just how important it is to prevent suspicion depends on the use and content of the multimedia data. Choosing a good background image should in most cases bring the level of suspicion to a reasonable level, although certain cases (such as the overlapping image object problem) require further measures to be taken. The entity creating the security objects should therefore 'preview' the image as requested by each available security classification in order to ensure that the level of suspicion is reduced to an acceptable level in every instance.

Any improvements to the region selection process should be welcomed since this should make region selection easier, faster, more accurate, and may even help to reduce suspicion. The basic security requirements as specified in this chapter for each phase of the model should always be adhered to in order for a system based on this model to maintain its purpose. Other security requirements are not as vital, and the enforcement of these depend on the sensitivity of the data and overall requirements of the computer system.