(12) UK Patent Application (19) GB (11) 2595679 (13) A

(43) Date of A Publication 08.12.2021

(21) Application No: 2008306.9

(22) Date of Filing: 02.06.2020

(71) Applicant(s):
Athlone Institute of Technology
Dublin Road, ATHLONE, County Westmeath, Ireland

(72) Inventor(s):
Nadia Kanwal
Mamoona Asghar
Marco Herbert
Brian Lee
Yuansong Qiao

(74) Agent and/or Address for Service:
Barker Brettell LLP
100 Hagley Road, Edgbaston, BIRMINGHAM,
B16 8QQ, United Kingdom

(51) INT CL:
G06T 7/194 (2017.01) ***G06F 16/535*** (2019.01)
G06F 16/735 (2019.01) ***G06T 9/00*** (2006.01)

(56) Documents Cited:
JP 2010278968 A **US 20020051491 A1**

(58) Field of Search:
INT CL **G06F, G06T, H04N**
Other: **WPI, EPODOC**

(54) Title of the Invention: **Video storage system**
Abstract Title: **Segmentation of an image into foreground and background, and subsequent encryption of both segments separately**

(57) Invention allows users to view security camera footage, whilst allowing personal data (i.e. image of person) to remain concealed from viewer. Video storage, comprising: segmentation of received video data frames into foreground 214 and background 215 data; separate encryption (Fig. 1; 104) of foreground and background data from each other; storing (Fig. 1; 106) encrypted foreground and background data separately. Segmentation may involve obtaining a reference frame (i.e. background of previous video frame), which is then subtracted from each frame of video. The reference frame may be updated periodically to account for changes in illumination (i.e. due to daylight). Encryption keys may be stored in second storage module (Fig. 1; 108). The system may comprise verification module (Fig. 4; 410) for calculating unique verification markers, which may comprise a hash, for each of foreground and background data. The system may comprise access module with three different levels of encryption access: first 612A shows encrypted foreground 613 and encrypted background 614 data; second 612B shows decrypted background data 616 and encrypted foreground data 613; third 612C shows decrypted foreground data 615 and decrypted background data 616. The system may allow for verification of the third clearance level to determine image tampering.

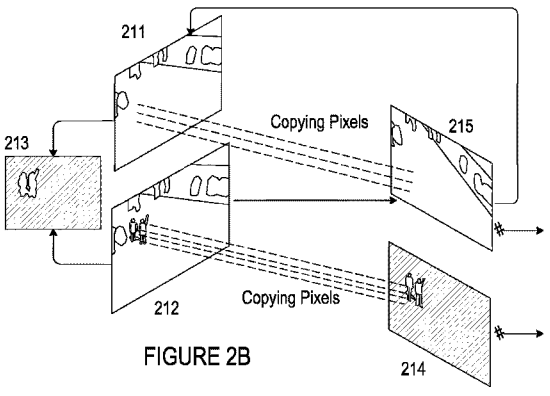

FIGURE 2B

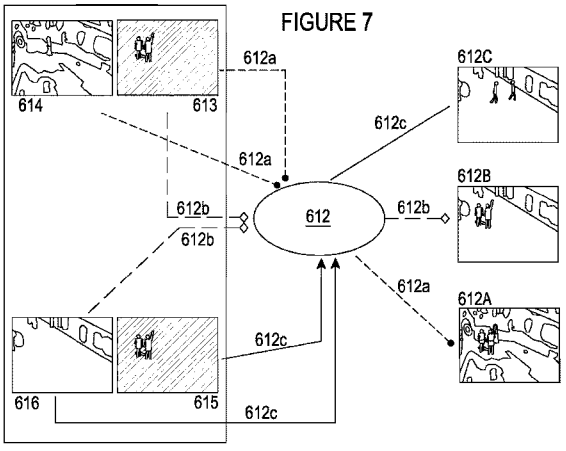

FIGURE 7

GB 2595679 A

At least one drawing originally filed was informal and the print reproduced here is taken from a later filed formal copy.

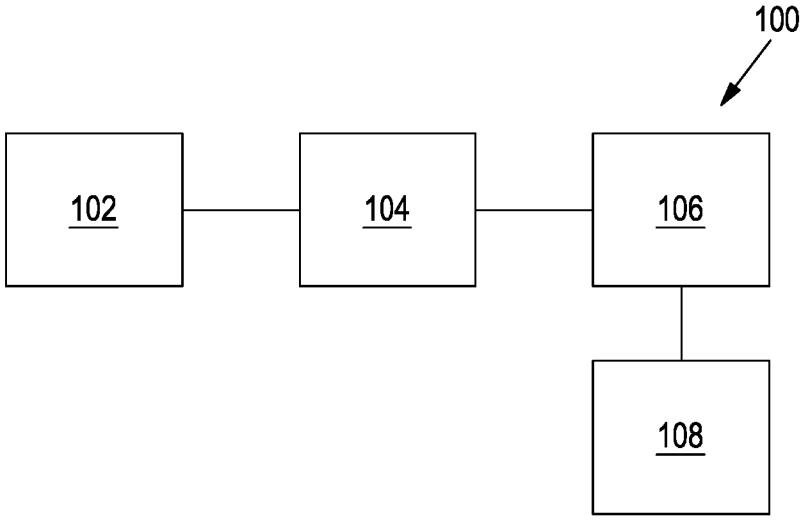

FIGURE 1

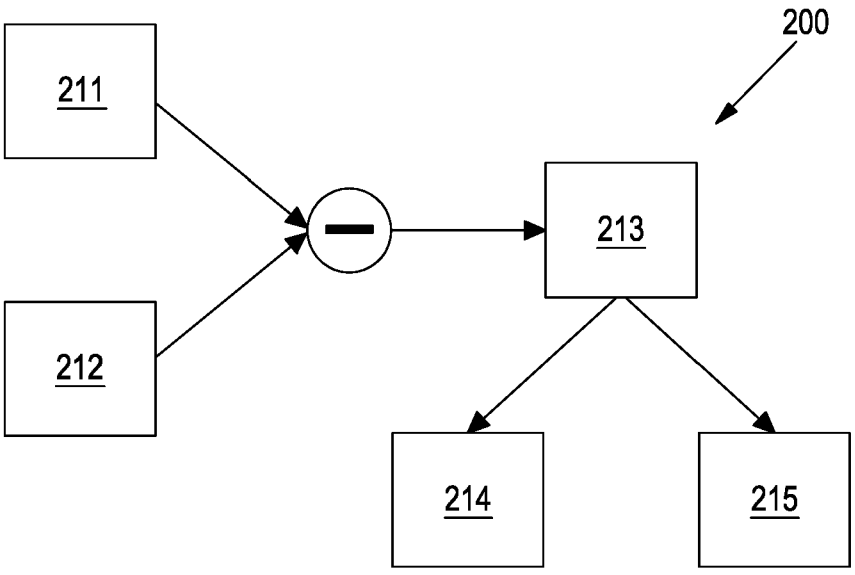

FIGURE 2A

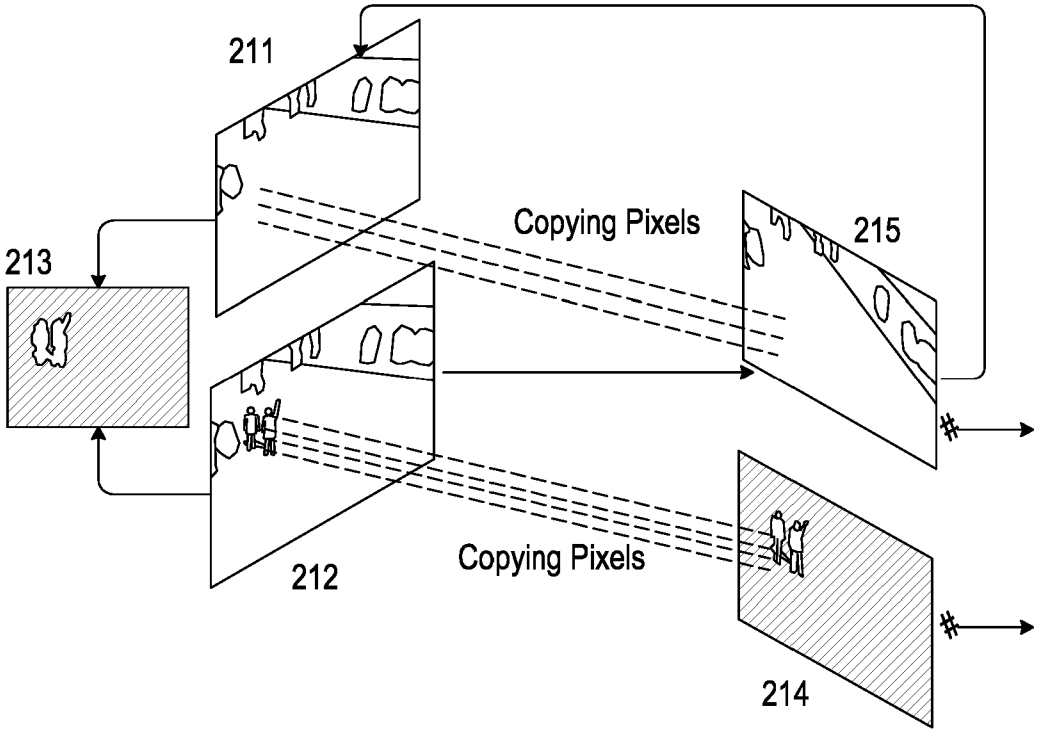

FIGURE 2B

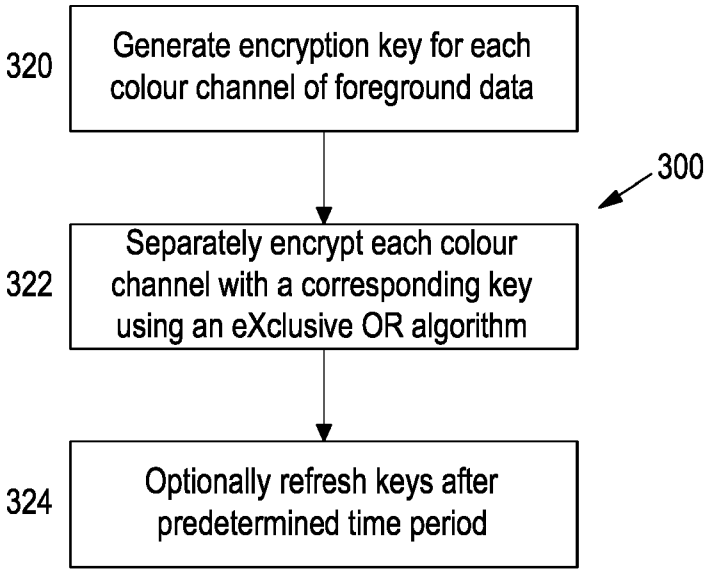

FIGURE 3A

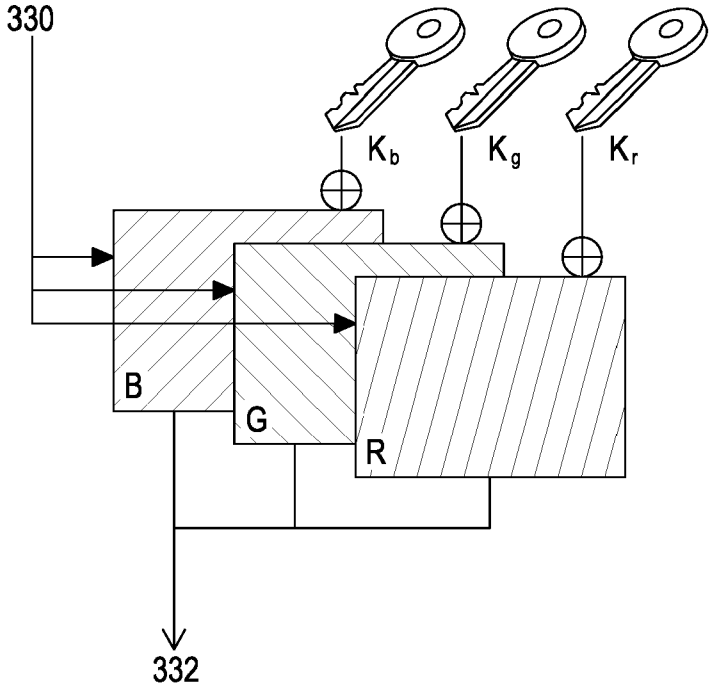

FIGURE 3B

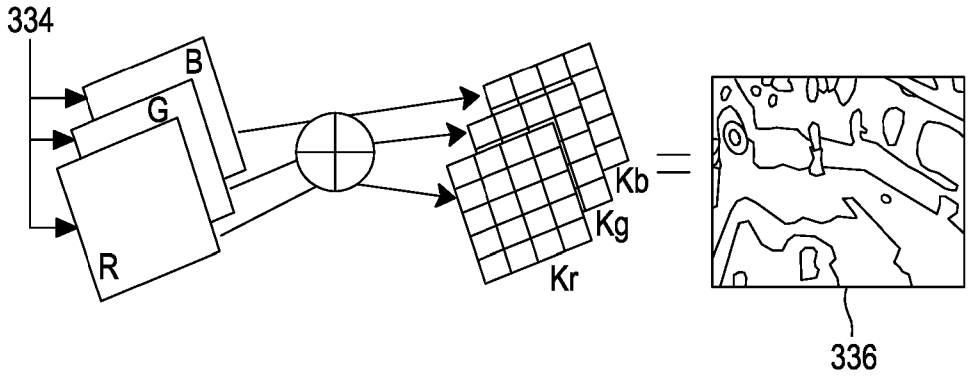

FIGURE 3C

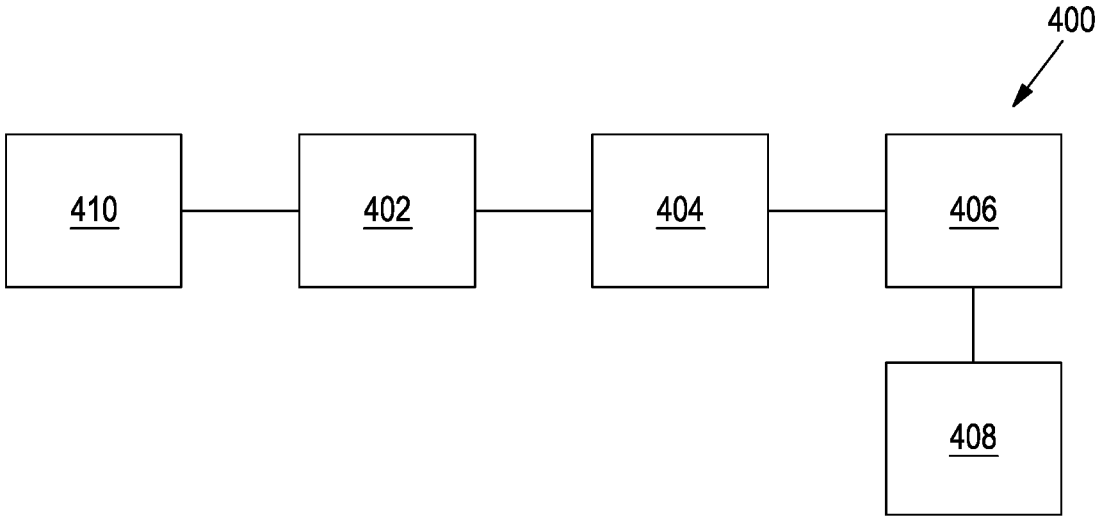

FIGURE 4

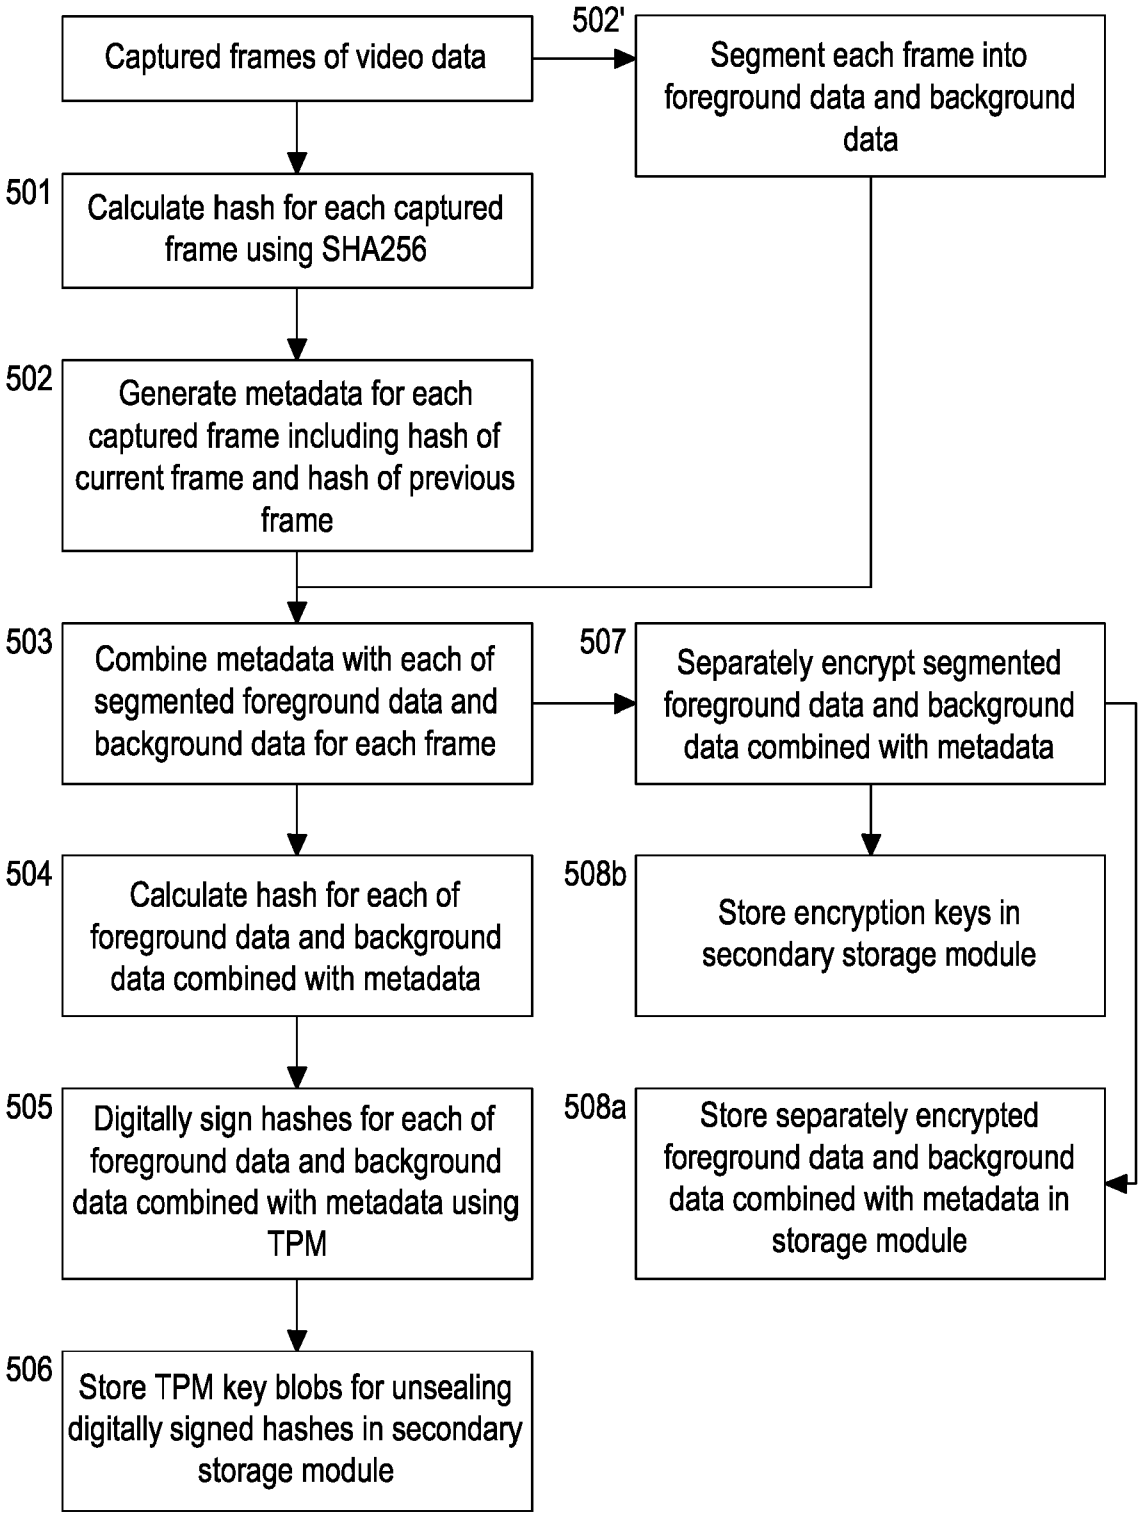

FIGURE 5A

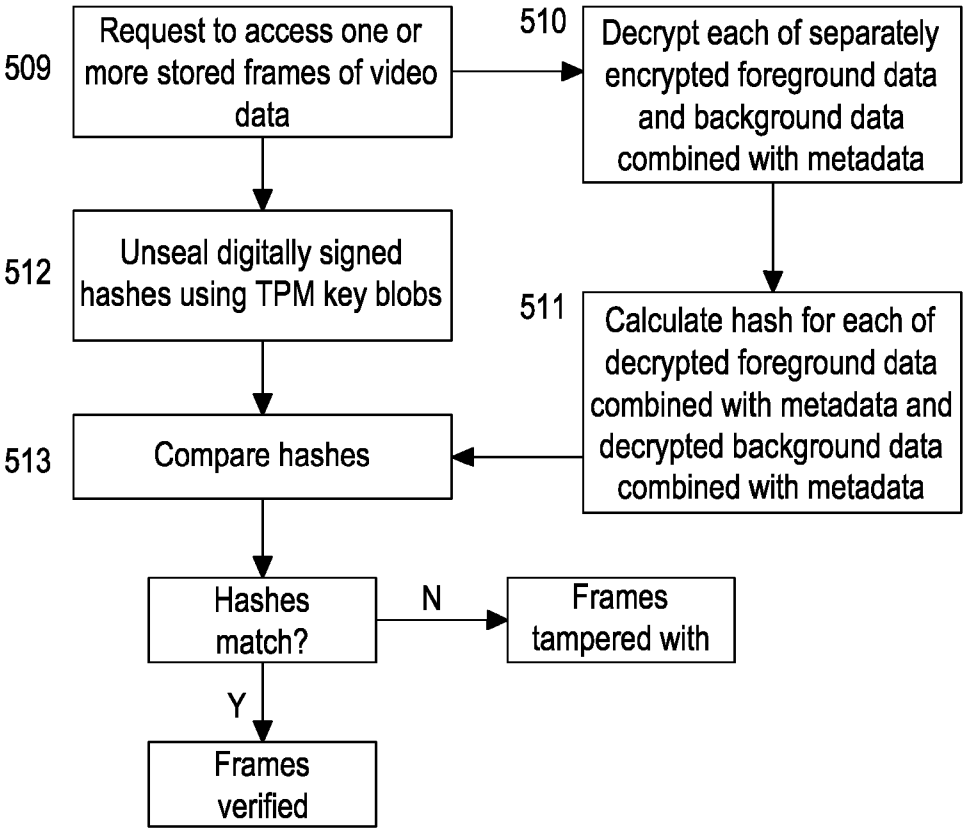

FIGURE 5B

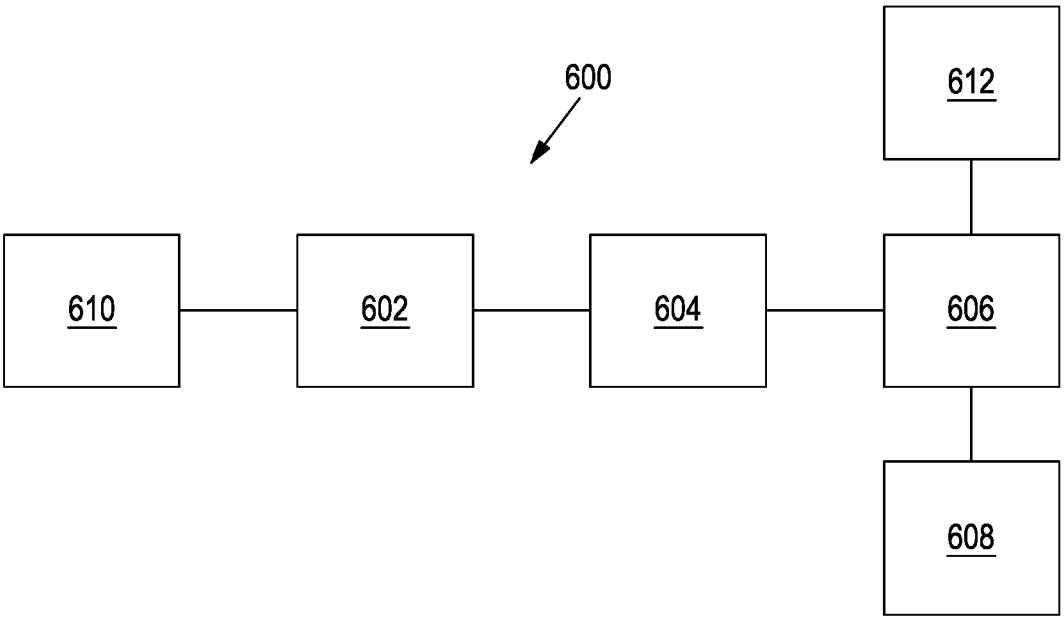

FIGURE 6

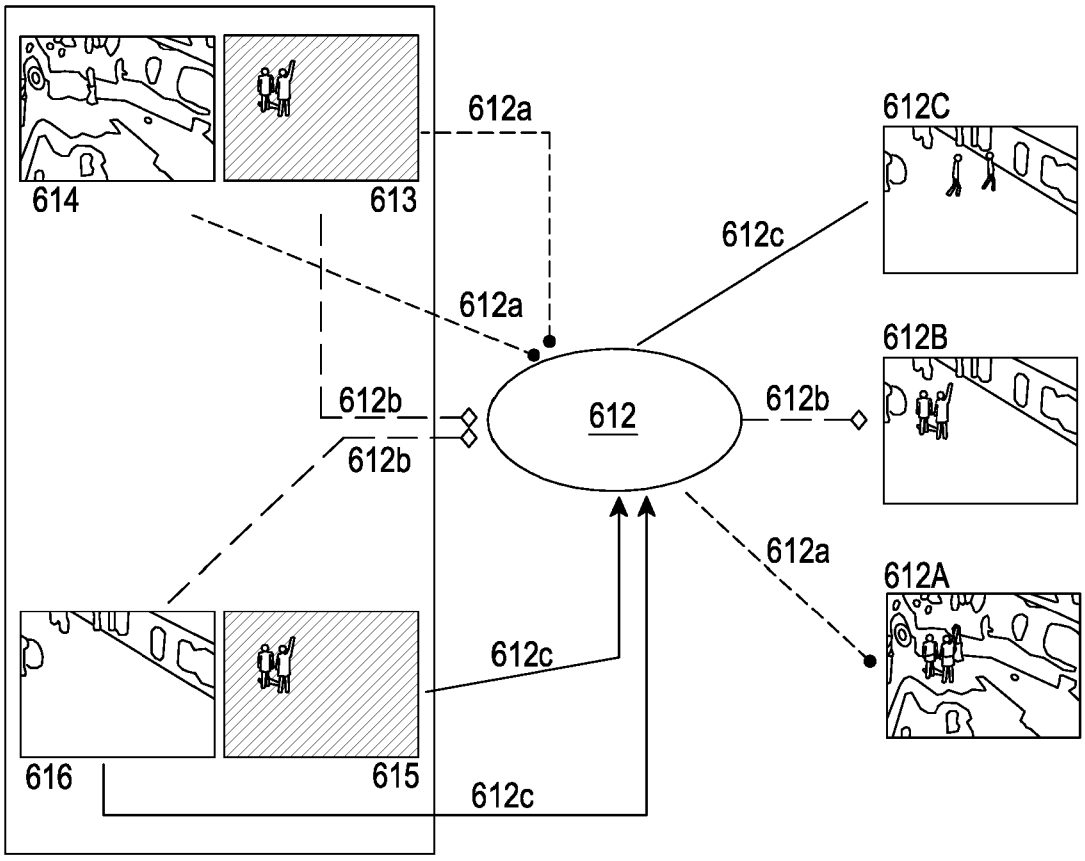

FIGURE 7

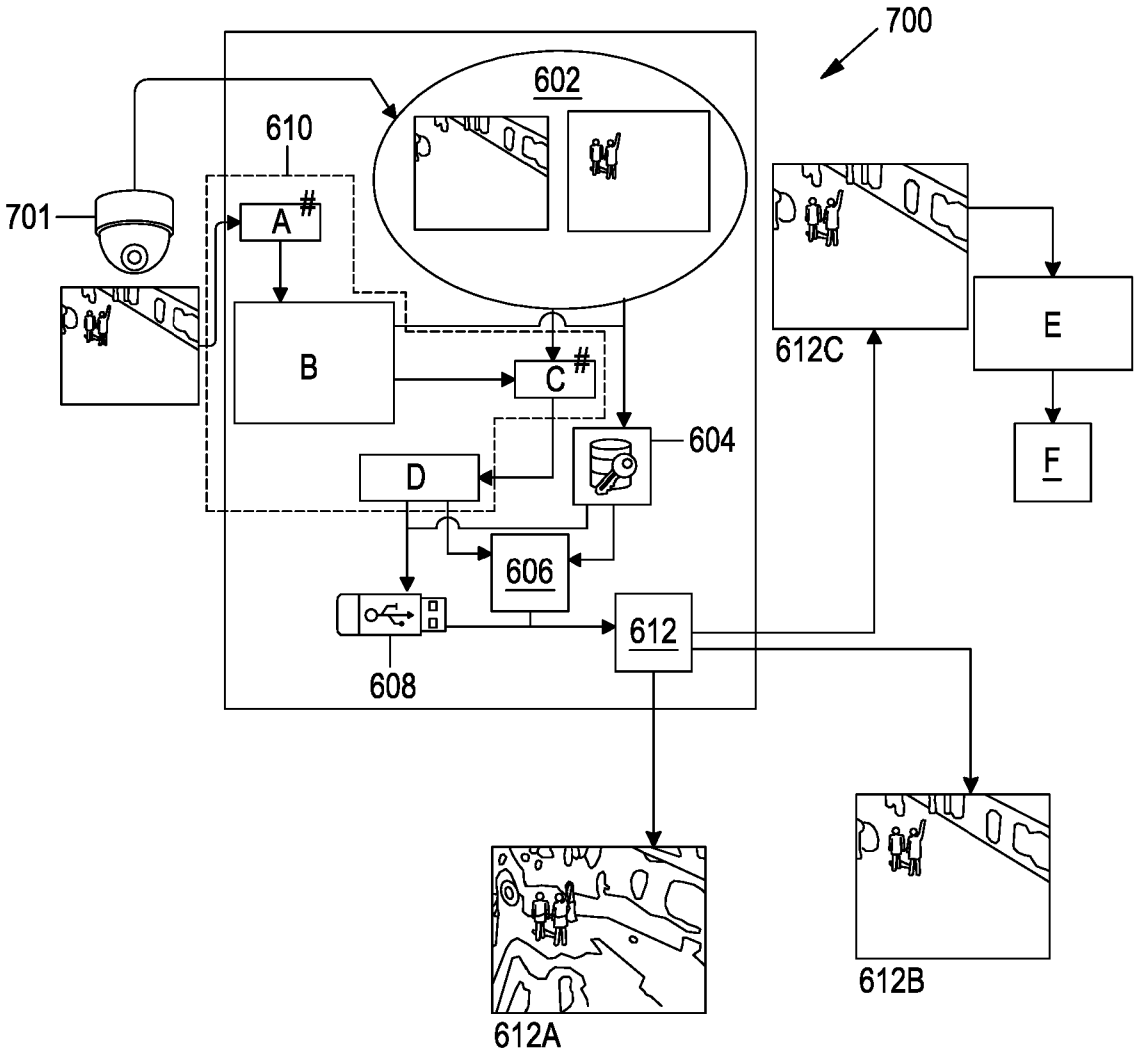

FIGURE 8

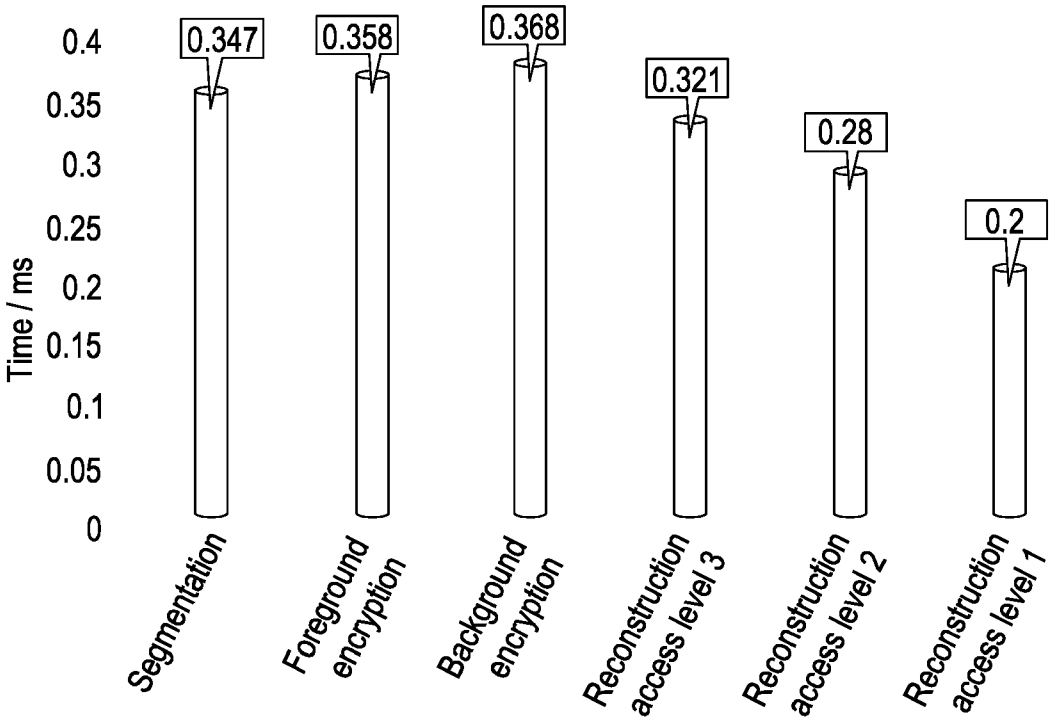

FIGURE 9

VIDEO STORAGE SYSTEM

ACKNOWLEDGEMENT

The projects leading to this application have received funding from Enterprise Ireland and the European Union's Horizon 2020 Research and Innovation Programme under the Marie Sklodowska-Curie grant agreement No 713654.

FIELD

The present invention relates to a video storage system and a video storage method.

BACKGROUND

Video imagery is generally considered an effective and efficient approach for carrying out surveillance. However, conventional video surveillance systems do not promise or provide privacy to either individuals or objects under surveillance. That may present both ethical problems and legal problems. Privacy protection may therefore hinder or prevent the installation and use of video surveillance systems, which may reduce security and increase a risk of damage or harm (for example, for items, locations or persons which require video surveillance).

The importance of personal data protection increased with the introduction of a new European Union (EU) regulation known as General Data Protection Regulation (GDPR) made effective in May 2018. That regulation also applies to visual data. In addition to privacy protection at the point of collection, many video surveillance systems also employ cloud-based storage facilities in the name of accessibility. According to GDPR, the use of cloud-based storage facilities threatens the privacy of stored visual data.

The present invention has been devised with the foregoing in mind.

SUMMARY

According to a first aspect, there is provided a video storage system. The video storage system may comprise a segmentation module. The segmentation module may be configured to receive video data. The segmentation module may be configured to receive video data from a camera in substantially real-time, or near real-time (for example, after a time delay). Alternatively, the segmentation module may be configured to receive pre-stored video data. The segmentation module may segment

each frame of the video data into foreground data and background data. The video storage system may also comprise an encryption module. The encryption module may be configured to encrypt the foreground data and background data of each frame of the video data separately from one another. The video storage system may further comprise a storage module. The storage module may be configured to store the encrypted foreground data and the encrypted background data of each frame of the video data separately from one another.

Segmenting each frame of the video data into foreground data and background data, for example on acquisition, and subsequently encrypting and storing the foreground data and the background data of each frame of the video data separately from one another may enable reversible privacy protection of foreground data without reducing effectiveness of a video surveillance system. The encrypted foreground data and the encrypted background data may be accessed, retrieved and optionally decrypted separately from one another. The retrieved and optionally decrypted foreground data and background data may be reconstructed to form a fully encrypted, partially encrypted or fully decrypted version of the original frame of video data. For example, decrypted background data may be combined with encrypted foreground data to reconstruct a partially encrypted frame of video data. That may enable video storage and retrieval which protects privacy by making foreground data (for example sensitive or targeted image content such as individuals and/or objects) non-identifiable, without inhibiting or preventing recognition or detection of activity (a key aspect of a video surveillance system). It may not be necessary to fully decrypt the frame of video data for the purpose of activity recognition or detection, so privacy protection may be maintained without sacrificing performance of a video surveillance system.

In contrast, conventional approaches encrypt video data captured from a camera as a whole before storing the encrypted video data, in order to provide privacy protection for stored video data. However, such conventional approaches clearly require the video data to be subsequently decrypted as a whole in order to obtain useful information from video surveillance. Privacy protection of individuals and/or objects is therefore not maintained upon retrieval of the video data from storage without post-processing techniques such as obfuscation or redaction to hide faces or other identifiable information. Post-processing of retrieved video data also results in additional time and complexity in order to maintain privacy protection after retrieval.

The system may be configured for use in a video surveillance system or network. The video surveillance system or network may comprise one or more cameras and one or more display monitors. The system may provide video storage and retrieval with privacy protection by default, without requiring additional post-processing after retrieval such as obfuscation or redaction. The segmentation module and the encryption module may together form a buffer for processing video data prior to storage and retrieval. The system may be used in a video surveillance system or network to enable captured video to be displayed or observed substantially live (for example, with a short delay of substantially 10 seconds or less allowing for segmentation, encryption, storage, retrieval and reconstruction) whilst still providing privacy protection of individuals and objects where necessary.

The encryption module may be configured to encrypt the foreground data and/or the background data of each frame of the video data using the same or a different encryption key or encryption key array (set of encryption keys). That may result in a large number of encryption keys being used to encrypt the video data, which may increase the strength of encryption of the video data.

The encryption module may be configured to encrypt each colour channel of each frame of the video data separately. The colour channels of the video data may comprise red, green and blue colour channels. The encryption module may be configured to encrypt each colour channel of each frame of the video data using the same or a different encryption key (e.g., using an encryption key array). A different encryption key array may be used to encrypt each colour channel of the foreground data of each frame of the video data and the background data of each frame of the video data respectively. Alternatively, the same encryption key or encryption key array may be used to encrypt each colour channel of the foreground data and the background data of each frame of the video data. That may further increase the strength of encryption of the video data.

The encryption module may be configured to encrypt each frame of the video data using an Exclusive OR (XOR) algorithm. An XOR algorithm may be used on constrained devices (for example, devices with limited or constrained energy, memory or processing capabilities) due to its inherently simple nature. Employing an inherently simple encryption algorithm such as an XOR encryption algorithm with a large number of encryption keys may enhance encryption strength whilst minimising processing requirements to encrypt the video data.

The encryption module may be configured to generate encryption keys to encrypt each frame of the video data. Alternatively, the encryption module may be configured to utilise pre-generated encryption keys to encrypt each frame of the video data. The system (for example, the encryption module) may be configured to regenerate one or more encryption keys after a pre-determined time period. Refreshing or regenerating encryption keys may protect the encryption keys from brute-force attacks.

The encryption module may be configured to encrypt each of the foreground data and the background data of each frame of the video data combined with metadata of the respective frame. The metadata may comprise a unique marker of the current frame and a unique marker of one or more preceding frames. That may provide a chain of unique markers which may enable a secure chain of evidence (for example, where video data may be used by law enforcement bodies or agencies).

The video storage system may further comprise a secondary storage module. The secondary storage module may be configured to store encryption keys generated by the encryption module. The secondary storage module may be removable from the video storage system. That may enable the encryption keys to be stored separately from the encrypted foreground data and encrypted background data, improving security. If the secondary storage module is physically removed from the system, the encryption keys may become inaccessible and the encrypted foreground data and encrypted background data may not be decrypted. The secondary storage module may employ two-factor authentication to enable stored data to be accessed. That may further improve security of stored data.

The video storage system may further comprise a verification module. The verification module may be configured to calculate a unique verification marker relating to each of the foreground data and the background data of each frame. The storage module may be configured to store the unique verification markers. The verification module may enable original captured video data to be reliably reconstructed from segmented foreground data and background data that has been encrypted and stored separately, without compromising privacy protection of the stored video data. The verification module may also enable a correct sequence or order of the frames of video data to be maintained.

The unique verification markers may be or comprise a hash of each of the foreground data and the background data of each frame. The unique verification markers may be or comprise a hash of each of the foreground data and the background

data of each frame combined with metadata of the frame. The metadata of the frame may comprise a unique marker of the frame and a unique marker of one or more preceding frames. That may provide a chain of unique markers which may enable a secure chain of evidence (for example, where video data may be used by law enforcement bodies or agencies). The unique markers may be or comprise, for example, a hash. The metadata may comprise one or more of a camera ID, a geographical location of the camera, a frame ID and a time stamp.

The verification module may be configured to digitally verify, e.g., digitally sign or seal the unique verification markers. The verification module may comprise a trusted platform module (TPM) configured to digitally sign or seal the unique verification markers. Digitally verifying/signing the unique verification markers may enable the video storage system to determine that the unique verification markers were generated by the video storage system. That may ensure that hackers do not attack the unique verification markers to make it appear that the video data has been tampered with. That provides a further level of security of the stored data.

The secondary storage module may be configured to store information configured to unseal the digitally signed unique verification markers. The secondary storage module may be removable from the video storage system. That may enable the information configured to unseal the digitally signed unique verification markers to be stored separately from the digitally signed unique verification markers, improving security. If the secondary storage module is physically removed from the system, the information configured to unseal the digitally signed unique verification markers may become inaccessible and the digitally signed unique verification markers may not be unsealed. The secondary storage module may employ two-factor authentication to enable stored data to be accessed. That may further improve security of stored data.

The video storage system may further comprise an access module. The access module may be configured to determine which of the encrypted foreground data and the encrypted background data of one or more requested frames of the video data should be decrypted. The access module may also be configured to reconstruct the requested frames based on the determination. That may enable privacy protection of the reconstructed video data to be maintained if the access module determines that a user is not authorized to see decrypted background data and/or decrypted foreground data.

The access module may be configured to provide one or a plurality of access roles. The access roles may each require the same or a different authorization and may

each provide a different reconstruction of the requested frames of video data. The access module may be configured to receive user input to determine which access role is associated with the request to access one or more frames of the video data. A plurality of access roles may enable the video storage system to provide multiple different views of the same data (for example, with differing levels of decryption) depending on the level of authorization.

The access module may be configured to instruct the encryption module to decrypt neither of the encrypted foreground data and the encrypted background data of one or more requested frames of video data. The access module may be configured to instruct the encryption module to decrypt neither of the encrypted foreground data and the encrypted background data of one or more requested frames of video data in response to the access module identifying a first access role input by a user. The first access role may correspond to a request for access to one or more frames of stored video data by an unauthorized user. That may prevent an unauthorized user from detecting or recognising activity or identifying objects and/or individuals from the reconstructed video data.

The access module may be configured to instruct the decryption module to decrypt only the encrypted background data of one or more requested frames of video data. The access module may be configured to instruct the decryption module to decrypt only the encrypted background data of one or more requested frames of video data in response to the access module identifying a second access role input by a user. The second access role may correspond to a request for access to one or more frames of stored video data by a video controller or video processor. That may enable activity to be detected or recognised from the reconstructed video data whilst maintaining privacy protection of objects and/or individuals forming part of the image content.

The access module may be configured to instruct the encryption module to decrypt both the encrypted foreground data and the encrypted background data of one or more requested frames of video data. The access module may be configured to instruct the encryption module to decrypt both the encrypted foreground data and the encrypted background data of one or more requested frames of video data in response to the access module identifying a third access role input by a user. The third access role may correspond to a request for access to one or more frames of stored video data by a law enforcement agency or body. That may enable original captured video data to be reconstructed for the purposes of recognising or detecting activity and identifying objects and/or individuals forming part of the image content.

When the access module identifies the third access role, the access module may be configured to instruct the verification module to calculate a unique verification marker relating to each of the decrypted foreground data and the decrypted background data of the one or more requested frames of video data. The access module may be configured to compare the calculated unique verification markers to the stored unique verification markers to determine whether the one or more requested frames of video data have been tampered with. That may enable the video storage system to verify that the reconstructed video data is original captured video data, in which case the reconstructed video data may be used by law enforcement bodies or agencies (for example, as evidence in criminal prosecution).

The access module may be configured to instruct the verification module to unseal the digitally signed unique verification markers. That may confirm that the unique verification markers were calculated by the video storage system, and may provide further confirmation that the reconstructed video data is original captured video data.

The access module may be configured to instruct the encryption module to encrypt the one or more reconstructed frames of video data before allowing a user access to the one or more reconstructed frames of video data (for example, before transferring or moving a copy of the one or more reconstructed frames out of the video storage system). That may preserve privacy protection and security of the video data stored by the video storage system, by preventing any of the original captured video data from leaving the video storage system in its original format.

The segmentation module may be configured to subtract a reference frame from each frame of video data to identify foreground data and background data. That may enable a simple approach to frame segmentation which has low processing requirements suitable for use on constrained devices.

The segmentation module may be configured to update the reference frame, for example to compensate for changes in illumination and/or changes in background image content (for example, one or more background objects being included within or removed from an image capture area of a camera). That may enable the frame segmentation to remain accurate over time, without requiring additional processing requirements or processing time. The reference frame may comprise the background data of the preceding frame of video data. The segmentation module may be configured to update the reference frame after a predetermined time period. Additionally or alternatively, the segmentation module may be configured to update

the reference frame on detecting a change in light intensity above a predetermined threshold.

According to a second aspect, there is provided a video surveillance system comprising the video storage system of the first aspect. The video surveillance system may comprise one or more cameras and one or more display monitors.

According to a third aspect, there is provided a method of storing video data. The method may comprise segmenting each frame of video data into foreground data and background data. The method may also comprise encrypting the foreground data and background data of each frame of the video data separately from one another. The method may also comprise storing the encrypted foreground data and the encrypted background data of each frame of the video data separately from one another.

Segmenting each frame of the video data into foreground data and background data on acquisition (for example, from a camera or using pre-stored video data), and subsequently encrypting and storing the foreground data and the background data of each frame of the video data separately from one another may enable reversible privacy protection of foreground data without reducing effectiveness of a video surveillance. The encrypted foreground data and the encrypted background data may be accessed, retrieved and optionally decrypted separately from one another. The retrieved and optionally decrypted foreground data and background data may be reconstructed to form a fully encrypted, partially encrypted or fully decrypted version of the original frame of video data. For example, decrypted background data may be combined with encrypted foreground data to reconstruct a partially encrypted frame of video data. That may enable video storage and retrieval which protects privacy by making foreground data (for example sensitive or targeted image content such as individuals and/or objects) non-identifiable, without inhibiting or preventing recognition or detection of activity (a key aspect of a video surveillance system). It may not be necessary to fully decrypt the frame of video data for the purpose of activity recognition or detection, so privacy protection may be maintained without sacrificing performance of video surveillance.

In contrast, conventional approaches encrypt video data captured from a camera as a whole before storing the encrypted video data, in order to provide privacy protection for stored video data. However, such conventional approaches clearly require the video data to be subsequently decrypted as a whole in order to obtain useful information from video surveillance. Privacy protection of individuals and/or objects is therefore not maintained upon retrieval of the video data from storage

without post-processing techniques such as obfuscation or redaction to hide faces or other identifiable information. Post-processing of retrieved video data also results in additional time and complexity in order to maintain privacy protection after retrieval.

The method may also be used in video surveillance to provide a buffer for processing video data prior to storage and retrieval. The method may enable captured video to be displayed or observed substantially live (for example, with a short delay of substantially 10 seconds or less allowing for segmentation, encryption, storage, retrieval and reconstruction) whilst still providing privacy protection of individuals and objects where necessary.

The method may comprise encrypting the foreground data and/or the background data of each frame of the video data using the same or a different encryption key or encryption key array.

The method may comprise encrypting each colour channel of each frame of the video data separately. The colour channels of the video data may comprise red, green and blue colour channels. The method may comprise encrypting each colour channel of each frame of the video data using the same or a different encryption key.

The method may comprise encrypting each frame of the video data using an Exclusive OR (XOR) algorithm.

The method may comprise regenerating one or more encryption keys after a predetermined time period.

The method may comprise encrypting each of the foreground data and the background data of each frame of the video data combined with metadata of the respective frame. The metadata may comprise a unique marker of the current frame and a unique marker of one or more preceding frames.

The method may comprise storing encryption keys in a secondary storage location separate from a storage location of the encrypted foreground data and encrypted background data. The secondary storage location may be removably attached to a system in which the storage location of the encrypted foreground data and encrypted background data is provided.

The method may comprise calculating a unique verification marker relating to each of the foreground data and the background data of each frame. The method may comprise calculating a unique verification marker of each of the foreground data and the background data of each frame combined with metadata of the frame. The unique verification markers may be a hash. The metadata of the frame may comprise a unique marker of the current frame and a unique marker of one or more preceding frames.

The unique marker may be or comprise, for example, a hash. The metadata may comprise one or more of a camera ID, a geographical location of the camera, a frame ID and a time stamp.

The method may comprise digitally verifying e.g. signing or sealing the unique verification markers. The method may comprise digitally signing or sealing the unique verification markers using a trusted platform module (TPM). The method may comprise storing information configured to unseal the digitally signed unique verification markers in a secondary storage location separate from a storage location of the digitally signed unique verification markers.

The method may further comprise accessing stored video data. The method may comprise determining which of the encrypted foreground data and the encrypted background data of one or more requested frames of the video data should be decrypted. The method may comprise reconstructing the requested frames based on the determination. The method may comprise displaying the reconstructed frames of video data.

The method may comprise determining which of the encrypted foreground data and the encrypted background data of the one or more requested frames should be decrypted based on an access role or authorization associated with a request to access one or more frames of the video data. The method may comprise a single access role, or may comprise a plurality of access roles. The method may comprise employing only a single access role of a plurality of available access roles.

The method may comprise decrypting neither of the encrypted foreground data and the encrypted background data of one or more requested frames of video data. The method may comprise decrypting neither of the encrypted foreground data and the encrypted background data of one or more requested frames of video data in response to identifying a first access role. The first access role may correspond to a request for access to one or more frames of stored video data by an unauthorized user.

The method may comprise decrypting only the encrypted background data of one or more requested frames of video data. The method may comprise decrypting only the encrypted background data of one or more requested frames of video data in response to identifying a second access role. The second access role may correspond to a request for access to one or more frames of stored video data by a video controller or video processor.

The method may comprise decrypting both the encrypted foreground data and the encrypted background data of one or more requested frames of video data. The

method may comprise decrypting both the encrypted foreground data and the encrypted background data of one or more requested frames of video data in response to identifying a third access role. The third access role may correspond to a request for access to one or more frames of stored video data by a law enforcement agency or body.

The method may comprise calculating a unique verification marker relating to each of the decrypted foreground data and the decrypted background data of the one or more requested frames of video data, for example in response to identifying the third access role. The method may further comprise comparing the calculated unique verification markers to unique verification markers calculated before encryption of the foreground data and the background data to determine whether the one or more requested frames of video data have been tampered with.

The method may comprise unsealing the digitally signed unique verification markers.

The method may comprise encrypting one or more reconstructed frames of video data before allowing a user access to the one or more reconstructed frames of video data (for example, transferring or moving a copy of the one or more reconstructed frames of video data).

The method may comprise subtracting a reference frame from each frame of video data to identify foreground data and background data. The method may comprise updating the reference frame, for example to compensate for changes in illumination and/or changes in background image content (for example, one or more background objects being included within or removed from an image capture area of a camera). The reference frame may comprise the background data of the preceding frame of video data. The method may comprise updating the reference background frame after a predetermined time period. Additionally or alternatively, the method may comprise updating the reference background frame on detecting a change in light intensity above a predetermined threshold.

According to a fourth aspect, there is provided a method of segmenting video data into foreground data and background data. The method may comprise subtracting a reference background frame from each frame of video data to identify foreground data and background data. That may provide a simple approach to frame segmentation which has low processing requirements suitable for use on constrained devices.

The method may comprise updating the reference frame, for example to compensate for changes in illumination and/or changes in background image content (for example, one or more background objects being included within or removed from an image capture area of a camera). That may enable the frame segmentation to remain accurate over time, without requiring additional processing requirements or processing time. The method may comprise updating the reference frame after a predetermined time period. Additionally or alternatively, the method may comprise updating the reference frame on detecting a change in light intensity above a predetermined threshold.

According to a fifth aspect, there is provided a method of encrypting video data. The method may comprise encrypting each frame of the video data separately using an Exclusive OR (XOR) algorithm. The method may comprise encrypting each frame of the video data using the same or a different encryption key or encryption key array.

That may enable an inherently simple XOR encryption algorithm to be employed with a large number of encryption keys, which may enhance encryption strength whilst minimising processing requirements to encrypt the video data. That approach may be particularly suitable for securely encrypting video data using constrained devices (for example, devices with limited or constrained energy, memory or processing capabilities).

The method may comprise encrypting each colour channel of the video data (for example, of each frame of the video data) separately. The colour channels may comprise red, green and blue colour channels. The method may comprise encrypting each colour channel of the video data (for example, of each frame of the video data) using the same or a different encryption key.

The method may comprise segmenting each frame of the video data into foreground data and background data. The method may comprise encrypting the foreground data and the background data separately from one another. The method may comprise encrypting the foreground data and the background data with the same or different encryption keys. That may increase the number of encryption keys used to encrypt the video data without substantially increasing processing requirements to encrypt the video data.

The method may comprise regenerating one or more encryption keys after a predetermined time period.

The method may comprise encrypting segmented foreground data and background data of each frame of the video data using an Exclusive OR (XOR)

algorithm. The method may comprise encrypting the foreground data and background data

Optional features of any of the above aspects may be combined with the features of any other aspect, in any combination. For example, features described in connection with the video storage system may have corresponding features definable with respect to the video surveillance system of the second aspect or the methods of the third, fourth or fifth aspects, and these embodiments are specifically envisaged. Features which are described in the context or separate aspects and embodiments of the invention may be used together and/or be interchangeable wherever possible. Similarly, where features are, for brevity, described in the context of a single embodiment, those features may also be provided separately or in any suitable sub-combination.

BRIEF DESCRIPTION OF DRAWINGS

The invention will now be described by way of example only with reference to the accompanying drawings in which:

FIG. 1 shows an embodiment of a video storage system in accordance with the invention;

FIGs. 2A and 2B show an embodiment of a frame segmentation method in accordance with the invention;

FIG. 3A shows an embodiment of an encryption method in accordance with the invention;

FIGs. 3B and 3C show a schematic of the encryption method shown in FIG. 3A;

FIG. 4 shows another embodiment of a video storage system in accordance with the invention, comprising a verification module;

FIGs. 5A and 5B show an embodiment of a data verification method in accordance with the invention;

FIG. 6 shows another embodiment of a video storage system in accordance with the invention, comprising an access module;

FIG. 7 shows an operational schematic of the access module of the video storage system of FIG. 6;

FIG. 8 shows an operational schematic of the video storage system of FIG. 6; and

FIG. 9 shows processing time for various operational steps carried out by the video storage system of FIG. 6.

Like reference numbers and designations in the various drawings may indicate like elements.

DETAILED DESCRIPTION

Figure 1 shows an embodiment of a video storage system 100. The video storage system 100 comprises a segmentation module 102. The video storage system 100 also comprises an encryption module 104. The video storage system 100 further comprises a storage module 106. An optional secondary storage module 108 is also shown, and will be discussed later.

The segmentation module 102 is configured to receive video data, e.g., from a camera and segment each frame of the video data into foreground data and background data. The purpose of segmentation is to separate sensitive or targeted image content from background content.

In the embodiment shown, the segmentation module 102 is configured to segment each frame of the video data into foreground data and background data using a frame subtraction method. An embodiment of a frame subtraction method 200 utilised by the segmentation module 102 is shown in more detail in Figure 2A.

Referring now to Figure 2A, a reference frame 211 is subtracted from a frame 212 of the video data resulting in a subtracted image 213. The subtracted image 213 is then separated into foreground pixels and background pixels. Each pixel in the subtracted image 213 having intensity above a predetermined threshold (for example, an intensity of greater than 20 in a range from 0 to 255) is classified as a foreground pixel. The remaining pixels in the subtracted image 213 having intensity below the predetermined threshold are classified as background pixels. Using a threshold to separate foreground pixels and background pixels facilitates noise removal and identifies areas of a video data frame 202 in which change of significance has occurred. Alternatively, a threshold may not be used, and any pixel having an intensity greater than 0 may be classified as a foreground pixel.

The foreground pixels of the subtracted image 213 indicate the locations of pixels of foreground data of the frame 212 whilst the background pixels of the subtracted image 213 indicate the locations of pixels of background data of the frame 212. Pixels in the frame 212 of the video data at locations corresponding to the

locations of foreground pixels in the subtracted image 213 are extracted or copied from the frame 212 to segment the foreground data 214 of the frame 212. The remaining pixels in the frame 212 of video data (at locations corresponding to the locations of background pixels in the subtracted image 213), in combination with pixels from the reference frame 211 at locations corresponding to the locations of foreground pixels in the subtracted image 213, form the segmented background data 215 of the frame 212. An example of the process described above is shown in Figure 2B.

In the embodiment shown in Figure 2B, an initial reference frame 211 is acquired at the initiation of a video surveillance session (for example, a session may be considered a period of time that the camera is continuously capturing video). The reference frame 211 is updated, e.g., to compensate for changes in illumination.

In the embodiment shown, the segmented background data 215 of the current frame 212 of video data is used as the reference frame 211 to segment the next frame 212 of video data into foreground data and background data, as shown in Figure 2B. In that way, the reference frame 211 is continuously updated to compensate for changes in illumination and improve segmentation performance. That approach may also enable verification of the original frame 212 of video data (discussed below).

Additionally or alternatively, the segmentation module 102 may be configured to update the reference frame 211 after a predetermined time period has elapsed, either from an initial reference frame being obtained or from a most recent update of the reference frame. The predetermined time period may be, for example, substantially one minute, five minutes, or 10 minutes, or 20 minutes, or one hour, but it will be appreciated that any suitable predetermined time period may be used. For example, for video data captured in an environment in which illumination is likely to vary significantly and frequently (for example external environments, due to weather changes, or due to shadows being cast by passing objects such as people or vehicles), the predetermined time period may be relatively short. Conversely, for video data captured in an environment in which illumination is likely to be substantially constant (for example internal environments with little exposure to natural light through windows), the predetermined time period may be relatively long.

Alternatively, or additionally, the segmentation module 102 may be configured to update the reference frame 211 on detection of a change in light intensity above a predetermined threshold. The segmentation module 102 may detect a change in light intensity by analysing pixel intensity in the frame 212. Alternatively, the segmentation

module 102 may be configured to detect a change in light intensity based on light intensity data collected at a location of the camera used to capture the video data (for example, data from a photodetector adjacent the camera used to capture the video data). The predetermined threshold may be a threshold relating to an absolute change in light intensity from the current reference frame 211, irrespective of illumination conditions for the current reference frame 211 (referred to as an absolute threshold). Alternatively, the predetermined threshold may be a threshold relating to a relative change in light intensity from the current reference frame 211, accounting for illumination conditions for the current reference frame 211 (referred to as a relative threshold). A relative threshold may enable the segmentation module 102 to update the reference frame to adapt to smaller changes in illumination for conditions of lower illumination than for conditions of higher illumination.

Alternatively, the reference frame 211 may not be updated to compensate for changes in illumination. The initial reference frame may be used continuously to segment each frame of video data into foreground data and background data for that video surveillance session. A new reference frame may be obtained at the initiation of each subsequent video surveillance session.

The segmentation module 102 may be configured to utilise other segmentation methods to segment each frame of video data into foreground data and background data. For example, the frame subtraction method described above may be modified to differentiate between humans and objects, for example using face recognition, gait recognition, or object detection. However, such modifications incur costs such as increased time and/or required processing power. The frame subtraction method described above may be suitable for use on constrained devices (for example, devices with limited or constrained energy, memory, processing speed etc.), as it may enable reduced processing requirements whilst still enabling segmentation of video data frames into foreground data and background data.

In addition, some object detection algorithms detect objects as image regions and therefore may include background data in the detected region, which may lead to background image content being wrongly classified as foreground data. That may prevent or inhibit location privacy, since wrongly classified background data may provide an indication of location.

Returning to Figure 1, the encryption module 104 is configured to encrypt the foreground data of each frame of the video data separately from the background data of each frame of the video data. The encryption module 104 is configured to receive

the segmented foreground data and background data of each frame from the segmentation module 102. The storage module 106 is configured to store the encrypted foreground data and the encrypted background data of each frame of the video data separately from one another.

An embodiment of an encryption method 300 utilised by the encryption module 104 is shown in more detail in Figure 3A. The method 300 is described with respect to foreground data of a frame of video data, but is equally applicable to background data of a frame of video data. A schematic of the method 300 is also shown in Figures 3B and 3C.

At step 320, an encryption key is generated (for example, by the encryption module 104) for each of the colour channels of the foreground data 330. In the embodiment shown, the colour channels comprise red R, blue B and green G colour channels resulting in three keys K_r, K_g, K_b being generated, although other colour channels (for example, having a different number of colour channels) may alternatively be used. In the embodiment shown, the keys K_r, K_g, K_b are symmetric keys dynamically generated at runtime generated by means of a pseudo-random function (PRF), but the keys may be generated in any suitable manner. At step 322, each colour channel R, G, B is separately encrypted by the corresponding generated key K_r, K_g, K_b using an eXclusive OR (XOR) algorithm, denoted by the symbol $\oplus$ in Figures 3B and 3C. In the embodiment shown, the foreground data 320 is encrypted by a key array comprising the generated keys K_r, K_g, K_b having a size of 24 bits (each generated key K_r, K_g, K_b having a size of 8 bits), such that each coloured pixel of the foreground data 330 is encrypted with the 24 bit key array. The method 300 therefore utilises a multi-channel XOR encryption approach to provide encrypted foreground data 332. At step 324, the keys K_r, K_g, K_b are refreshed after a predetermined time period. In the embodiment shown, the predetermined time period is 24 hours, although the keys may be refreshed after any suitable time period has elapsed. For example, depending on the location of the camera used to capture the video data, the keys may be refreshed after a time period of anything up to substantially one week. Refreshing the keys may protect the keys from brute-force attacks. Alternatively, the keys may not be refreshed.

Figure 3C shows an example of an encrypted frame 336 of video data encrypted according to the method 300 described above. A segmented frame 334 of video data is provided to the encryption module 104. The foreground data and the background data of the frame 334 are encrypted separately from one another, with

each colour channel R, G, B of both the foreground data and the background data encrypted separately as described above. In the embodiment shown, the same set of encryption keys K_r, K_g, K_b is used to encrypt the colour channels R, G, B of both the foreground data and the background data of the frame 334. As can be seen in Figure 3C, identification of individuals, objects and activity in the encrypted frame 336 is not possible due to the encryption. Alternatively, a different set of encryption keys may be used to encrypt the respective colour channels R, G, B of the foreground data and the background data of the frame 334.

The simplicity of XOR approaches means that XOR approaches have not been widely implemented in modern encryption schemes. However, for constrained devices, the simplicity of XOR approaches provides an advantage. By encrypting each colour channel of each of the foreground data and background data of each frame of video data separately with a different encryption key, the data may be securely encrypted whilst minimising processing requirements. For example, a video of 60 seconds in length and having 30 frames per second (fps) is made up of 1800 frames. Segmenting each frame of the video data into foreground data and background data, and then further dividing each of the foreground data and the background data into red, green and blue colour channels, could result in 10,800 different encryption keys for 60 seconds of video data, if a different set of encryption keys is used for each of the red, green and blue colour channels of each of the foreground data and the background data of each frame. Even if the same set of encryption keys is used for each colour channel of both the foreground data and the background data of each frame, 5,400 different encryption keys could be used for 60 seconds of video data at 30 fps. That increases the difficulty for hackers to attack and decrypt the video data. A multi-channel XOR encryption approach as described above may therefore be suitable for use on constrained devices, as it may minimise processing requirements (due to use of an inherently simple encryption algorithm) whilst still enabling secure encryption of the video data.

Alternatively, the encryption module 104 may not encrypt each colour channel of the foreground data and the background data of each frame separately (or each frame of the video data may comprise a grayscale image). However, the encryption module 104 may still employ XOR encryption using a different encryption key for each of the foreground data and the background data of each frame. That provides the advantage of being able to use numerous XOR encryption keys simultaneously for each frame of the video data to increase the strength of XOR encryption whilst

minimising processing requirements to encrypt the video data. Alternatively, the encryption module 104 may employ XOR encryption but may use the same encryption key for each of the foreground data and the background data of each frame (with a different encryption key being used for each frame).

Alternatively, the encryption module 104 may be configured to encrypt each of the foreground data and the background data of each frame separately from one another using any suitable encryption scheme. For example, an industry standard encryption scheme such as the Advanced Encryption Standard (AES) may be used to encrypt each of the foreground data and the background data of each frame separately from one another. The encryption module 104 may also be configured to encrypt each colour channel of each of the foreground data and the background data of each frame separately from one another using any suitable encryption scheme.

The encryption module 104 may be configured to use a key management scheme to enhance security of keys generated to encrypt the foreground data and the background data of each frame (and optionally to encrypt each colour channel of the foreground data and the background data of each frame). A key management scheme may incur computational and networking costs, and so may not be implemented on a constrained device.

Once the foreground data and the background data have been encrypted separately from another, the encrypted foreground data and the encrypted background data are stored separately from one another using the storage module 106. In the embodiment shown, the storage module 106 is or comprises an edge device. However, the stored data may be easily accessed through an IP address, provided correct authorization is provided. Alternatively, the storage module 106 may be or comprise cloud-based storage.

The encrypted foreground data and the encrypted background data can be retrieved from the storage module 106 separately from one another. The separate encryption and storage of the foreground data and the background data may also enable the foreground data and background data to be decrypted separately from one another. For example, neither, one or both of the encrypted foreground data and the encrypted background data may be encrypted. That may enable a fully encrypted, partially encrypted (for example, only background data decrypted) or fully decrypted version of the original frame of video data to be reconstructed from the retrieved and (optionally) decrypted foreground data and background data.

In the embodiment shown in Figure 1, the system 100 comprises a secondary storage module 108. The secondary storage module 108 is configured to store encryption keys generated by the encryption module 104. The encrypted video data (encrypted foreground data and encrypted background data of each frame of the video data) is stored on the storage module 106. In the embodiment shown, the secondary storage module 108 is a secure external storage device (such as a hardware wallet, commonly used to store cryptocurrencies) that is removably attached to the system 100. In the embodiment shown, the secondary storage module 108 employs two-factor authentication in order to access data (such as encryption keys generated by the encryption module) stored on the secondary storage module 108. Any suitable two-factor authentication means or technique may be utilised. Therefore, if the secondary storage module 108 is physically removed from the system 100, the encryption keys on the secondary storage module 108 are inaccessible and the encrypted video data stored on the storage module 106 cannot be decrypted.

Alternatively, the secondary storage module 108 may not be removably attached to the system 100, but may be fixedly or permanently attached to the system 100. Encryption keys generated by the encryption module 104 may still be stored on the secondary storage module 108, separate from the encrypted video data stored on the storage module 106. Storing the encryption keys and the encrypted video data separately in physically different locations in the system 100 may still provide enhanced security of the encrypted video data. Alternatively, the encryption keys generated by the encryption module 100 may be stored in the storage module 106 together with the encrypted video data.

Figure 4 shows another embodiment of a video storage system 400. The video storage system 400 comprises a segmentation module 402, an encryption module 404, a storage module 406 and a secondary storage module 408 as described above with respect to the video storage system 100.

The video storage system 400 also comprises a verification module 410. The verification module 410 may ensure data integrity of video data stored by the video storage system 400. The verification module 410 may enable original captured video data to be reliably reconstructed from segmented foreground data and background data (that has been separately encrypted and stored), without compromising privacy protection of the stored video data. Reconstructed and verified original captured video data may be used, for example, by law enforcement bodies or agencies.

Figures 5A and 5B show an embodiment of a data verification method 500 which can be utilised by the video storage system 400 comprising the verification module 410. Figure 5A shows how frames of video data are prepared for data verification. Figure 5B shows how stored frames of video data are verified following a retrieval request.

Starting with Figure 5A, at step 501 a unique verification marker is calculated for each captured frame of video data by the verification module 410. In the embodiment shown, the unique verification marker is a hash calculated using SHA256 (a member of the Secure Hash Algorithm 2 family), but the unique verification markers may be calculated using any suitable approach.

At step 502, metadata is generated for each captured frame of video data. In the embodiment shown, the metadata is generated by the verification module 410, but it will be appreciated that metadata may be generated by another part of the video storage system 400. In the embodiment shown, the metadata of a frame comprises a unique marker (for example a hash) of the frame (the current frame) and a unique marker (for example a hash) of the previous frame. Alternatively, the metadata of a frame may comprise a unique marker of the current frame and a unique marker of a plurality of previous frames of video data captured by the same camera. For example, the metadata of a frame may comprise a unique marker for each of a number of most immediately previous frames (for example, two, three, four, five or more frames directly preceding the current frame), or may comprise a unique marker for each previous frame captured by the same camera. Linking of the unique markers of one or more preceding frames of video data with the unique marker of the current frame provides a chain of unique markers for captured frames, which may maintain a secure chain of evidence (discussed below). Additionally or alternatively, the metadata of a frame may comprise one or more of a camera ID, a geographical location of the camera (for example, GPS coordinates), and a frame ID (for example comprising a serial number, date and time of the frame).

In the embodiment shown, steps 501 and 502 described above take place separately (for example, prior to or in parallel with) to segmentation of the current video frame into foreground data and background data. Segmentation of the video frame into foreground data and background data takes place at step 502'. The segmentation module 402 may be configured to segment the video frame into foreground data and background data as described above with respect to the segmentation module 102.

At step 503, the verification module 410 combines (for example, attaches or otherwise associates) the metadata of the frame (generated at step 502 described above) with each of the foreground data and the background data of the frame separately.

At step 504, the verification module 410 calculates a unique verification marker for each of the foreground data combined with the metadata of the frame and the background data combined with the metadata of the frame. In the embodiment shown, the unique verification marker is a hash calculated using SHA256 (a member of Secure Hash Algorithm 2 family), but the unique verification marker may be calculated using any suitable approach.

Alternatively, metadata may not be generated for each captured frame of video data. A unique verification marker (for example a hash) may be calculated for each of the foreground data and the background data alone, without being combined with metadata of the frame.

At step 505, the verification module 410 digitally signs (or seals) the unique verification markers relating to each of the foreground data and the background data. In the embodiment shown, the unique verification markers are calculated for each of the foreground data and the background data combined with metadata of the frame, but it will be appreciated that step 505 is equally applicable to unique verification markers calculated for each of the foreground data and the background data without metadata. In the embodiment shown, the verification module 410 comprises a trusted platform module (TPM) configured to digitally sign the unique verification markers. In the embodiment shown, the TPM employs public-private key pairs generated using a Rivest-Shamir-Adleman (RSA) encryption algorithm. The TPM seed is generated using a pseudo-random function (PRF) and is induced to generate further keys by using a root key of the TPM. To digitally sign the unique verification markers relating to each of the foreground data and the background data, the TPM signs the unique verification markers with a private key of a public-private key pair generated by the TPM. The digitally signed unique verification markers are stored on the storage module 406. Alternatively, other suitable approaches for digitally signing the unique verification markers may be used.

At step 506, the video storage system 400 stores information configured to unseal the digitally signed unique verification markers. In the embodiment shown, the information configured to unseal the digitally signed unique verification markers is stored on the secondary storage module 408, but may be alternatively be stored on the

storage module 406 along with the digitally signed unique verification markers. Evidently the information configured to unseal the digitally signed unique verification markers needs to be stored in an accessible location in order to unseal the digitally signed unique verification markers. In the embodiment shown, the information configured to unseal the digitally signed unique verification markers comprises a public key corresponding to the private key of the generated public-private key pair used to digitally sign the unique verification markers. The public key corresponding to the private key of the generated public-private key pair can be used to decrypt (or unseal) the unique verification markers to verify that the unique verification markers were in fact calculated by the video storage system 400 (discussed further below). In the embodiment shown, the public key is contained in a TPM key blob encrypted using a key hierarchy of the TPM. The TPM key blob is stored on the secondary storage module 408 in the embodiment shown, but may alternatively be stored on the storage module 406. Alternatively, the information configured to unseal the digitally signed unique verification markers may depend on the approach used to digitally sign the unique verification markers.

Alternatively, the unique verification markers may not be digitally signed, and may instead be stored in as-calculated form on the storage module 406.

At step 507, the video storage system 400 separately encrypts the segmented foreground data and background data using the encryption module 404. In the embodiment shown, encryption module 404 encrypts each of the segmented foreground data and background data combined with metadata of the frame. Alternatively, the segmented foreground data and background data may be encrypted without being combined with metadata of the frame. The encryption module 404 may separately encrypt the segmented foreground data and background data as described above with respect to the encryption module 104.

At step 508a, the separately encrypted foreground data and background data are stored in the storage module 406 as described above with respect to the storage module 106. At step 508b, encryption keys used to separately encrypt the foreground data and background data are stored in the secondary storage module 408 as describe above with respect to the secondary storage module 108. Alternatively, the encryption keys may be stored on the storage module 406.

Steps 501 to 508 are shown in Figure 5A and illustrate one embodiment of how frames of video data can be prepared for data verification.

Turning to Figure 5B, at step 509 the video storage system 400 receives a request to access one or more frames of video data stored on the storage module 406.

At step 510, the encryption module 404 decrypts both the segmented foreground data and background data of each of the one or more requested frames of video data retrieved from the storage module 406. In the embodiment shown, the segmented foreground data and background data of each of the one or more requested frames of video data are encrypted in combination with metadata of the frame, but that is not essential (as described above). In the embodiment shown, the encryption module 404 decrypts the segmented foreground data and background data using encryption keys retrieved from the secondary storage module 408. Alternatively, the encryption keys may be stored on and retrieved from the storage module 406.

At step 511, the verification module 410 calculates a unique verification marker relating to each of the decrypted foreground data and the decrypted background data of each of the one or more requested frames of video data. In the embodiment shown, the unique verification marker is a hash calculated using SHA256 as described above. However, the verification module 410 may use any suitable approach to calculate unique verification markers, as long as the same approach to calculate unique verification markers is used throughout the data verification method 500.

At step 512, the digitally signed unique verification markers generated at step 505 are retrieved from the storage module 406 and unsealed by the verification module 410 using the information configured to unseal the digitally signed verification markers retrieved from the secondary storage module 408. As described above, in the embodiment shown the information configured to unseal the digitally signed unique verification markers comprises a public key corresponding to a private key of a generated public-private key pair used to digitally sign the unique verification markers. The public key is contained in a TPM key blob stored on the secondary storage module 408. Alternatively, the information configured to unseal the digitally signed unique verification markers may be stored on and retrieved from the storage module 406.

At step 513, the unique verification markers calculated at step 511 are compared with the unique verification markers unsealed at step 512. If the unique verification markers match one another, the one or more requested frames of video data are considered verified. Matching unique verification markers (for example hashes) provides an indication that the one or more requested frames have not been

tampered with, because the verification markers generated for each of the segmented foreground data and background data are unique to the data. If the unique verification markers match, then the segmented foreground data and background data can be combined to reconstruct original captured video data. The verification module 410 employing a data verification method such as the data verification method 500 described above may ensure that the reconstructed video is identical to original captured video data and not an altered version of the original captured video data.

The data verification method 500 described above requires that both the separately encrypted foreground data and background data of each frame be decrypted in order to verify authenticity of the video data. The verification module 410 may be used to verify only partially decrypted video data (for example, frames of video data where only background data is decrypted to preserve privacy of individuals or objects in the foreground data). That is because unique verification markers are calculated for each of the foreground data and the background data of each frame separately from one another. However, because the authenticity of the still encrypted foreground data cannot be verified without decrypting it, the authenticity of the whole of the video data cannot be verified. The verification module 410 may therefore provide most benefit when both the segmented foreground data and background data are decrypted when attempting to reconstruct original captured video data.

Figure 6 shows another embodiment of a video storage system 600. The video storage system 600 comprises a segmentation module 602, an encryption module 604, a storage module 606 and a secondary storage module 608 as described above with respect to the video storage systems 100, 400. The video storage system also comprises a verification module 610 as described above with respect to the video storage system 400.

The video storage system 600 further comprises an access module 612. The access module 612 is configured to determine which of the separately encrypted foreground data and background data of each frame should be decrypted in response to a request for access to one or more frames of video data stored by the video storage system 600. The access module 612 is configured to do so using a hierarchical authorization structure comprising a plurality of access levels. The access module 612 is also configured to reconstruct a fully encrypted, partially encrypted (for example, only background data decrypted) or fully decrypted version of each frame of video data, depending on an access level identified. Figure 7 shows an operational schematic of an embodiment of the access module 612.

In the embodiment shown in Figure 7, the access module 612 is configured to use three access levels or access roles. A first access level 612a is provided for unauthorized users, as indicated by the dotted lines. If the access module 612 identifies that an unauthorized user has requested access to one or more frames of video data stored by the storage module 606 of the video storage system 600, the access module 612 determines that neither of the encrypted foreground data and the encrypted background data should be decrypted, and instructs the encryption module 604 accordingly. The access module 612 therefore retrieves encrypted foreground data 613 and encrypted background data 614. For the first access level 612a, a reconstructed frame 612A of video data output by the access module 612 is shown with both encrypted foreground data 613 and encrypted background data 614, providing a fully protected view of the frame 612A of video data. In a fully protected view of the frame 612A, neither identification of individuals and/or objects nor activity recognition or detection are possible.

In the embodiment shown, a second access level 612b is provided by the access module 612 for individuals employed as video controllers or video processors, as indicated by the dashed lines. If the access module 612 identifies that an individual authorized according to the second access level 612b has requested access to one or more frames of video data stored by the storage module 606 of the video storage system 600, the access module 612 determines that the encrypted background data should be decrypted but that the encrypted foreground data should remain encrypted, and instructs the encryption module 604 accordingly. In the embodiment shown, the encrypted background data is decrypted using an XOR encryption key array used to encrypt the data originally, as described above with respect to the encryption module 104. The encryption key array is retrieved from the secondary storage module 608, as described above with respect to the secondary storage module 108. However, any suitable encryption scheme may be used to encrypt and decrypt the background data. The access module therefore retrieves encrypted foreground data 613 and decrypted background data 616. For the second access level 612b, a reconstructed frame 612B of video data output by the access module 612 is shown with encrypted foreground data 613 and decrypted background data 616, providing a privacy protected view of the frame 612B of video data. The second access role 612b may therefore enable video surveillance capable of detecting or recognising activity whilst maintaining privacy of objects and/or individuals forming part of the image content.

In the embodiment shown, the access module 612 also provides a third access level 612c for use by law enforcement bodies or agencies. If the access module 612 identifies that an individual authorized according to the third access level 612c has requested access to one or more frames of video data stored by the storage module 606 of the video storage system 600, the access module 612 determines that both the encrypted foreground data and the encrypted background data should be decrypted, and instructs the encryption module 604 accordingly. The access module 612 therefore retrieves decrypted foreground data 615 and decrypted background data 616. For the third access level 612c, a reconstructed frame 612C of video data output by the access module 612 is shown with decrypted foreground data 615 and decrypted background data 616, providing a fully decrypted view of the frame 612C of video data.

The reconstructed frame 612C of video data output by the access module 612 may only be of use to law enforcement agencies or bodies (for example, for the purposes of criminal prosecution) if it can be verified that the reconstructed frame 612C is the same as the original frame of video data captured by a camera. In the embodiment shown, the access module 612 is configured, after retrieving decrypted foreground data 615 and decrypted background data 616, to instruct the verification module 610 to calculate a unique verification marker for each of the decrypted foreground data 615 and the decrypted background data 616. The verification module 610 may calculate the unique verification markers as described above with respect to the verification module 410. In the embodiment shown, for the third access level 612c, the access module 612 is further configured to retrieve stored unique verification markers relating to each of the foreground data and the background data (calculated before encryption and storage of the foreground data and background data by the encryption module 604 and the storage module 606 respectively) from the storage module 606. The access module 612 is also configured to compare the calculated unique verification markers with the stored unique verification markers, as described above with respect to the method 500, in order to determine whether the frame 612C is original captured video data, or the frame 612C represents video data that has been tampered with (as described above with respect to the method 500). If the stored unique verification markers have been digitally signed prior to storage on the storage module 606 (for example, using a TPM as described above with respect to the verification module 410), the access module 612 may instruct the verification module 610 to unseal the digitally signed unique verification markers prior to comparison.

If the compared unique verification markers match one another, the access module 612 determines that the frame 612C is original captured video data. In the embodiment shown, the access module 612 instructs the encryption module 604 to encrypt the reconstructed frame 612C prior to transfer or copying of the reconstructed frame 612C to an external storage device for use by the law enforcement body or agency (for example, as evidence in criminal prosecution). That is done in order to preserve protection of the video data stored by the video storage system 600, by preventing any of the original captured video data from leaving the video storage system 600 in its original format. The corresponding encryption key or keys used to encrypt the reconstructed frame 612C are also provided (for example, copied or transferred to the same external storage device) to the law enforcement body or agency. The encryption module 604 may encrypt the reconstructed frame using any suitable encryption scheme. For example, an XOR encryption scheme may be employed to encrypt the reconstructed frame 612C, or alternatively an industry standard encryption scheme such as AES may be used to encrypt the reconstructed frame 612C. Alternatively, the reconstructed frame 612C may not be encrypted prior to transfer or copying of the reconstructed frame 612C to an external storage device.

It will be appreciated that the access module 612 as described above may be utilised in a video storage system that does not comprise a verification module such as the verification modules 410, 610 described above, for example the video storage system 100 described above. The access module 612 may still provide a hierarchical authorization structure comprising a plurality of access levels configured to provide one or more reconstructed frames of video data at differing levels of accessibility or decryption. The reconstructed frame(s) of video data may not be verified as original captured video data. However, for scenarios other than use of video data as evidence in criminal prosecution, data verification may not be essential.

Alternatively, the access module 612 may provide a plurality of access levels (for example, two, three, four or more) each requiring a different authorization and providing a different reconstruction of the frames of video data. For example, the access module 612 may be configured to provide a reconstruction of a frame of video data comprising decrypted foreground data and encrypted background data, for an appropriate authorization. Additional access levels may be added to the access module 612 to accommodate specific reconstruction requirements.

In video storage systems not comprising an access module 612 (such as video storage systems 100, 400 described above), for example for use in conventional video

surveillance systems of private property, a default setting may be to provide reconstructed frames of video data in accordance with the second access level 612b. The video storage system may therefore automatically provide privacy protection of objects and/or individuals forming part of the image content whilst enabling detection or recognition of activity to maintain functionality of a video surveillance system. The video storage system may still comprise a verification module (such as verification modules 410, 610 described above) to enable verification of reconstructed original captured video data comprising decrypted foreground data and decrypted background data, if required.

Figure 8 shows an operational schematic of the video storage system 600 in its entirety as part of a video surveillance system or network 700. The video surveillance system 700 comprises at least one camera 701. The camera 701 captures video data. Each frame of video data captured by the camera 701 is segmented into foreground data and background data by the segmentation module 602, as shown by the separate images representing foreground data (right hand image comprising two individuals on a black background) and background data (left hand image). The verification module 610 (shown by the area enclosed by dashed lines in Figure 8) first calculates a unique marker for each frame of the video data captured by the camera (step A, as described above with respect to the method 500). In the embodiment shown, the unique verification marker comprises a hash calculated using SHA256. The verification module 610 then generates metadata for each frame comprising at least unique marker of the current frame, a unique marker of the previous frame (step B, as described above with respect to the method 500). Optionally the metadata generated for each frame comprises one or more of a camera ID, a geographical location of the camera (for example, GPS coordinates), and a frame ID of the current frame (for example comprising a serial number, date and time of the frame). The verification module 410 then calculates a unique verification marker for each of the segmented foreground data and the segmented background data combined with the metadata of the frame (step C, as described above with respect to the method 500). The unique verification markers are then digitally signed by the verification module (step D, as described above with respect to the method 500). The encryption module 604 encrypts each of the segmented foreground data and background data combined with the metadata of the frame. The separately encrypted foreground data and background data are stored on the storage module 606, whilst the encryption keys are stored on the secondary storage module 608. Similarly, the digitally signed unique verification markers are stored on

the storage module 606, whilst the information configured to unseal the digitally signed unique verification markers (for example, TPM key blobs) is stored on the secondary storage module 608.

When the access module 612 receives a request to access one or more frames of video data, the access module 612 identifies which of the three access levels 612a, 612b, 612c the request is being made through, and determines which of the encrypted foreground data and the encrypted background data should be decrypted. The access module 612 then retrieves the encrypted foreground data and the encrypted background data from the storage module 606, retrieves the encryption keys from the secondary storage module 608, and instructs the encryption module 604 accordingly (to decrypt neither or both of the encrypted foreground data and encrypted background data, or to decrypt only the encrypted background data). The access module 612 then reconstructs the frames of video data using the decrypted or encrypted foreground data and background data appropriate for each access level. The reconstructed frames 612A, 612B and 612C are formed for the first, second and third access levels 612a, 612b, 612c as described above. The access module 612 is configured to display the reconstructed frame according to the relevant access level on a display monitor (not shown) of the video surveillance system 700. For the reconstructed frame 612C formed for the third access level 612c, the access module 612 then verifies the reconstructed frame 612C is original captured video data or represents video data that has been tampered with (step E, as described above with respect to the method 500 and the video storage system 600). If the reconstructed frame 612C is determined to be original captured video data, the access module 612 instructs the encryption module 604 to encrypt the reconstructed frame 612C before transferring or copying the reconstructed frame 612C to an storage device external to the video storage system 600 (step F, as described above with respect to the video storage system 600).

An example system showing feasibility of the video storage system 600 is now described. The example system comprised a testbed comprising TPM enabled Intel NUC Core i5-6260U minicomputer, with Ubuntu 16.04, 64-bit OS and device specifications including 4 GB RAM and 1.80 GHz processor. The evaluation video data used was the publicly available dataset named 'Atrium' (available from https://www.jpjodoin.com/urbantracker/dataset.html). The evaluation video is 2.43 MB in size with Super Graphics Video Array (SVGA) resolution (800 x 600 pixels per frame) at 30 frames per second (fps).

For performance evaluation, processing time and video quality metrics were calculated. The processing time of the video storage system 600 for various stages is shown in Figure 9. The segmentation, encryption and reconstruction performed by the segmentation module 602, the encryption module 604 and the access module 612 did not take significant time. With a camera speed of 30 fps, the video storage system 600 takes approximately 10 seconds to segment, encrypt and store the evaluation data on the storage module 606. Reconstruction of the video data involved retrieving the encrypted data stored on the storage module 606, decrypting the encrypted foreground data and the encrypted background data as appropriate using the encryption module 604, and combining the resulting foreground data and background data using the access module. As can be seen from Figure 9, segmentation of the video data took 0.347 ms, encryption of the foreground data took 0.358 ms, and encryption of the background data took 0.368 ms. Figure 9 also shows that reconstruction according to the third access 612c (where both encrypted foreground data and encrypted background data needs to be decrypted) took 0.321 ms, reconstruction according to the second access level 612b took 0.28 ms (where only encrypted background data is decrypted) and reconstruction according to the first access level 612a took 0.2 ms (where neither encrypted foreground data nor encrypted background data needs to be decrypted, thereby consuming no time on decryption).

Table 1 (below) shows quality metrics for encrypted, decrypted and reconstructed frames of video obtained using the example video storage system 600. Video degradation through a multi-channel XOR encryption approach, as described above with respect to the method 300, is estimated using objective video quality metrics. The objective video quality metrics comprise: peak signal-to-noise ratio (PSNR) (for the luma (Y) and two UV chroma components to check one or more of the channels is still not secure even after encryption); Mean Squared Error (MSE); and Structural Similarity Index Metric (SSIM), matching human perception more closely than either PSNR or MSE. The values shown in the second row of Table 1 indicate a significant video quality degradation introduced by the privacy-protection of the video storage system 600. The values clearly indicate that the video quality for the first access level 612a is significantly poorer than for the second and third access levels 612b, 612c. The values in the bottom row therefore indicate that the video content has been completely protected such that an unauthorized user, having no access to encryption keys, cannot easily take advantage of the video for detection or identification purposes.

Access Level	PSNR{Y, U, V} dB	SSIM	MSE
Level 3	{33.0, 36.2, 36.8}	0.967	23.89
Level 2	{27.6, 30.1, 29.9}	0.948	83.28
Level 1	{17.5, 15.5, 12.3}	0.738	859.59

Table 1 Video quality metrics for each of the access levels provided by video storage system 600

From reading the present disclosure, other variations and modifications will be apparent to the skilled person. Such variations and modifications may involve equivalent and other features which are already known in the art of video surveillance and storage, and which may be used instead of, or in addition to, features already described herein.

Although the appended claims are directed to particular combinations of features, it should be understood that the scope of the disclosure of the present invention also includes any novel feature or any novel combination of features disclosed herein either explicitly or implicitly or any generalisation thereof, whether or not it relates to the same invention as presently claimed in any claim and whether or not it mitigates any or all of the same technical problems as does the present invention.

Features which are described in the context of separate embodiments may also be provided in combination in a single embodiment. Conversely, various features which are, for brevity, described in the context of a single embodiment, may also be provided separately or in any suitable sub-combination. The applicant hereby gives notice that new claims may be formulated to such features and/or combinations of such features during the prosecution of the present application or of any further application derived therefrom. Features of the devices and systems described may be incorporated into/used in corresponding methods. Where features are disclosed in connection with one embodiment of a video storage system, it should be appreciated that any one or more or all of the same features may be incorporated in other embodiments of video storage systems, instead of or in addition to the features described for the particular embodiment. That is, any and all combinations of features are envisaged, and are envisaged to be interchangeable, replaceable, added or removed. In particular, systems and methods comprising features relating to segmentation, encryption, data verification and data retrieval are specifically

envisaged to be provided either alone or in any combination. For example, a system or method may comprise features relating to segmentation and encryption, segmentation and encryption and data retrieval, segmentation and encryption and data verification and data retrieval, encryption and data verification and data retrieval, etc.

For the sake of completeness, it is also stated that the term "comprising" does not exclude other elements or steps, the term "a" or "an" does not exclude a plurality, a single processor or other unit may fulfil the functions of several means recited in the claims and any reference signs in the claims shall not be construed as limiting the scope of the claims.

CLAIMS

1. A video storage system comprising:

a segmentation module configured to receive video data and segment each frame of the video data into foreground data and background data;

an encryption module configured to encrypt the foreground data and background data of each frame of the video data separately from one another;

a storage module configured to store the encrypted foreground data and the encrypted background data of each frame of the video data separately from one another.

2. The video storage system of claim 1, wherein the encryption module is configured to encrypt the foreground data and the background data of each frame of the video data using a different encryption key or encryption key array.

3. The video storage system of claim 1 or of claim 2, wherein the encryption module is configured to encrypt each colour channel of each frame of the video data separately, and optionally wherein the encryption module is configured to encrypt each colour channel of each frame of the video data using a different encryption key.

4. The video storage system of any preceding claim, wherein the encryption module is configured to encrypt each frame of the video data using an Exclusive OR algorithm.

5. The video storage system of any preceding claim, wherein the system is configured to regenerate one or more encryption keys after a pre-determined time period.

6. The video storage system of any preceding claim, wherein the encryption module is configured to encrypt each of the foreground data and the background data of each frame combined with metadata of the frame.

7. The video storage system of any preceding claim, wherein the system further comprises a secondary storage module configured to store encryption keys generated

by the encryption module, and optionally wherein the secondary storage module is removable from the system.

8. The video storage system of any preceding claim, wherein the system further comprises:

a verification module configured to calculate a unique verification marker relating to each of the foreground data and the background data of each frame; and, optionally,

wherein the storage module is configured to store the unique verification markers.

9. The video storage system of claim 8, wherein the verification module is configured to calculate a unique verification marker for each of the foreground data and the background data of each frame combined with metadata of the frame, and optionally wherein the unique verification markers each comprise a hash.

10. The video storage system of claim 6 or of claim 9, wherein the metadata of the frame comprises a unique marker of the frame and a unique marker of one or more previous frames, and optionally wherein the unique markers each comprise a hash.

11. The video storage system of claim 10, wherein the metadata of the frame further comprises one or more of a camera ID, a geographical location of the camera, a frame ID and a time stamp.

12. The video storage system of any of claims 8 to 11, wherein the verification module is configured to digitally sign the unique verification markers.

13. The video storage system of claim 12, wherein the verification module comprises a trusted platform module configured to digitally sign the unique verification markers.

14. The video storage system of claim 12 or of claim 13 as dependent from claim 7, wherein the secondary storage module is configured to store information configured to unseal the digitally signed verification markers.

15. The video storage system of any preceding claim, wherein the system further comprises an access module configured to:

determine which of the encrypted foreground data and the encrypted background data of one or more requested frames of the video data should be decrypted; and

reconstruct the requested frames based on the determination.

16. The video storage system of claim 15, wherein the access module is configured to instruct the encryption module to:

i) decrypt neither of the encrypted foreground data and the encrypted background data of one or more requested frames of video data in response to the access module identifying a first access role input by a user; or

ii) decrypt only the encrypted background data of one or more requested frames of video data in response to the access module identifying a second access role input by a user; or

iii) decrypt both the encrypted foreground data and the encrypted background data of one or more requested frames of video data in response to the access module identifying a third access role input by a user.

17. The video storage system of claim 16, part iii), as dependent directly or indirectly from claim 8, wherein when the access module identifies the third access role, the access module is configured to:

instruct the verification module to calculate a unique verification marker relating to each of the decrypted foreground data and the decrypted background data of the one or more requested frames of video data;

compare the calculated unique verification markers to the stored unique verification markers to determine whether the one or more requested frames of video data have been tampered with.

18. The video storage system of claim 17 as dependent directly or indirectly from claim 12, wherein the access module is configured to instruct the verification module to unseal the digitally signed unique verification markers.

19. The video storage system of claim 17 or of claim 18, wherein the access module is configured to instruct the encryption module to encrypt the one or more

reconstructed frames of video data before allowing a user access to the one or more reconstructed frames of video data.

20. The video storage system of any preceding claim, wherein the segmentation module is configured to subtract a reference frame from each frame of video data to identify foreground data and background data.

21. The video storage system of claim 20, wherein the segmentation module is configured to update the reference frame to compensate for changes in illumination.

22. The video storage system of claim 21, wherein the reference frame comprises the background data of the preceding frame of video data.

23. The video storage system of claim 21, wherein the segmentation module is configured to update the reference frame:

i) after a predetermined time period; and/or

ii) on detecting a change in light intensity above a predetermined threshold.

Intellectual Property Office

Application No:	GB2008306.9	**Examiner:**	Steve Williams
Claims searched:	1 - 23	**Date of search:**	4 November 2020

Patents Act 1977: Search Report under Section 17

Documents considered to be relevant:

Category	Relevant to claims	Identity of document and passage or figure of particular relevance
A	--	JP 2010278968 A (PANASONIC)
A	--	US 2002/0051491 A1 (PHILLIPS ELECTRONICS)

Categories:

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

Field of Search:

Search of GB, EP, WO & US patent documents classified in the following areas of the UKC^{X} :

Worldwide search of patent documents classified in the following areas of the IPC

G06F; G06T; H04N

The following online and other databases have been used in the preparation of this search report

WPI, EPODOC

International Classification:

Subclass	Subgroup	Valid From
G06T	0007/194	01/01/2017
G06F	0016/535	01/01/2019
G06F	0016/735	01/01/2019
G06T	0009/00	01/01/2006