Blockchain Enabled Integrity Protection for Bodycam Video

Michael Kerr(✉), Fengling Han, and Ron Van Schyndel

School of Science (Computer Science), RMIT University, Melbourne, Australia
{michael.kerr,fengling.han,ron.vanschyndel}@rmit.edu.au

Abstract. The prevalence of both documented incidents and anecdotal evidence perpetuate mistrust in video collected via Law Enforcement body worn recording devices. This paper examines the application of blockchain technology for the management of high volumes of video produced every day during the course of a police field officers' duties. We apply a comprehensive blockchain system developed specifically for law enforcement video collection to the body worn scenario and examine the protection level offered whilst considering the specific requirements and limitations of this mobile platform. Specific scenarios are examined and shown to offer a compelling level of assurance to mobile body worn video collection operations.

Keywords: Law enforcement · Bodycam · Video · Mobile · Digital watermarking · Blockchain

1 Introduction

Due to the rapid expansion of the use of body worn cameras (BWC, Bodycams) by Law Enforcement Agencies (LEA) worldwide [1,2] citizen groups and individuals alike have voiced concern over the proliferation of discrete, mobile and government controlled surveillance technology [2,3]. And rightly so, communities have the right to expect that their public infrastructure is there to service the people's best interests, in this context being public safety and security. It is not always immediately apparent that this is the case, which is unfortunate as this is undoubtedly the original intention of such technology. In order to assure community faith and comfort in the ubiquitous presence of mobile LEA captured video records there must be transparency in the collection, storage, retrieval and use of this data [4]. There is a strong and urgent requirement that collected video records are not only genuine, good faith representation of policing events, but that their collection is also a matter of public record; in addition to integrity checking mechanisms there must also be a record of creation, and in some legislative areas also a record of destruction. Goold [2] makes the profound argument that whilst body worn cameras share many similarities with existing deployed state sanctioned surveillance technology, they further encroach on individuals, as they are by definition mobile, and not only recording in public spaces, but

W. Xiang et al. (Eds.): BROADNETS 2021, LNICST 413, pp. 163–173, 2022.
https://doi.org/10.1007/978-3-030-93479-8_10

private locations like personal residences if the officer is called to such locations. One peak civil liberty group, the American Civil Liberty Union (ACLU) holds the sanguine position that body worn cameras have the potential to protect all involved parties, if only the integrity of the collected product can be assured [5]. The challenge of maintaining tight access control whilst also facilitating public accountability motivates the implementation discussed in this paper. The systems employed to meet these requirements must keep pace with the surveillance technologies themselves, and must be transparent and ubiquitous, so that in time communities can consider the audit functions as an equally integral component of the public safety and police accountability effort these body worn cameras are applied to.

1.1 Background

The clear requirement to protect the integrity of LEA collected footage is a good fit for Distributed Ledger Technology (DLT, or Blockchain) and there are currently several such research projects orientated towards CCTV in the public space, and of those some specifically intersect with to field of Bodycam video and blockchain technology. A common approach is to utilise public networks such as the ERC20 smart contract capability of the Ethereum network through 3rd party suppliers that leverage this network [6]. Our own earlier work outlined the utility of adapting existing audit frameworks within law enforcement procedure [7] and implemented these frameworks using complementary technologies of Blockchain ledgers and digital watermarking implemented on the camera itself [8]. This infrastructure applies itself well to not just general LEA surveillance operations, but specifically the challenges surrounding bulk collection of body worn video by officers in the field.

Whilst leveraging existing public blockchain networks is a valid strategy, advantages of our system include:

- Providing a self-contained Merkle-tree based blockchain system that has no reliance on the continued existence of any financially driven blockchain network.
- Being an entirely independent system that can be implemented at any required security classification or network.
- We further enhance the system by facilitating the on-camera creation of blocks.

Our system is applied here to protect the integrity of Bodycam video, as well as its associated metadata in a distributed fashion that is autonomous from public DLT networks, and can be distributed within LEA, governance authorities or independent third parties. Critically, it is also decoupled from the video data itself.

1.2 Bodycam Use Cases and Challenges

There are a multitude of developed Bodycam products in use worldwide. Mobile capture of video is a core feature, and devices can record locally in a range of

resolutions, with some offering real time streaming over cellular and integration with local sensor networks, such as vehicle reed switches or Land Marine Radios (LMR). Scores of Bodycam devices collecting video every day generate vast amounts of data that is required to be centrally archived by the LEA. Different vendors approach this workflow with some variation of offloading data from the device at end of shift with the unit in a docking cradle or connected to a stations Wi-Fi network. Whilst some products can live stream when configured to do so, this is usually a tactical on demand function. Generally devices can be considered to be collecting video data in an independent manner and uploading centrally when a low cost, high bandwidth network becomes available.

Applicable Mechanisms from Camera to Client. Applying identification and integrity measures such as digital hashing, digital watermarking, or Distinctive DC Sequence operations on the camera itself can enable integrity protection at the earliest possible stage of the data's collection. Due to the partially offline, independent workflow of Bodycams in addition to the actual capture of video, there is increased requirement to *a.* Record the capture as an event, *b.* record video integrity information, *c.* Record any relevant metadata. We focus on how to process generated video and the immutable recording of this metadata, without impacting the device's image recording performance. Differing integrity protection mechanisms have their own practical implications to being deployed on the camera:

- Digital hashing mechanisms such as SHA or MD5 are well established and used heavily throughout the law enforcement and legal professions. They are well understood, and their output is widely accepted. Unfortunately hashing has significant limitations in our video scenario. Due to its binary output, in the event of a hash mismatch there is no information on where or how the data has altered from its original state, rendering the entire file suspect from an evidentiary perspective. Bodycams can be exposed to extreme environmental or tactical conditions, as well, this places particular risk on storage hardware and consequently hashing as a protective mechanism, even when deployed at the granular "key frame" level as is done in some video management systems.
- Digital Watermarking overcomes many of the issues surrounding the binary output of hashing. Many watermark schemes offer location based integrity information within the video down to inter frame location. Utilising watermarks trades the definitive nature of hashing for a metric capable of providing a level of confidence as well as information on what and where data could be modified, delivered within the above described volatile environment. This comes at a computational cost and it can be difficult to implement complex algorithms in real time on small, embedded hardware. Additionally, many vendors choose to deliver their own specific compression algorithms, complicating transform domain watermarking support across multiple vendors. This suggests that it is more practical to implement real time on-camera watermarking in the spatial domain as has been found by others specifically researching bench-marking [9]. After meeting implementation challenges it

needs to be also considered that, unsurprisingly, many CCTV vendors consider image quality a primary metric and differentiator to their products. This creates an imperative to minimise the introduction of noise into the video, further limiting the options of what watermark implementations are usable in a practical sense. For these reasons' visible watermarks, such as QR or bar codes can prove useful due to their known impact on the image, their ease of embedding and decoding and their open source.

In this system we utilise a visible watermark to link the video to a blockchain record, and within that record we store the sequence of DC values (position (0,0)) from each Discrete Cosine Transform (DCT) block within each triggered frame. At any subsequent point any user in possession of an enrolled video clip can query the blockchain and obtain metadata such as the distinctive DC Sequence (DCS) to examine the integrity of the video, without requiring access to the original device, video archive or any specifically protected metadata.

As a low complexity, multi-threaded operation that can be processed outside the multimedia pipeline of video capture our method of DCS enrolment is an integrity protection measure that can be implemented on small, embedded hardware, in real time and across many vendors, as it operates prior to, and independently of, the compression operation. Figure 1 shows the concurrent processing and creation of the blockchain data outside the multimedia pipeline of a Gstreamer implementation [10]. An attractive aspect for law enforcement is that, apart from the visible marking, this process in no way alters the main body of the frame and introduces no noise to the resulting video. It can therefore be presented in legal proceedings as unaltered with that caveat.

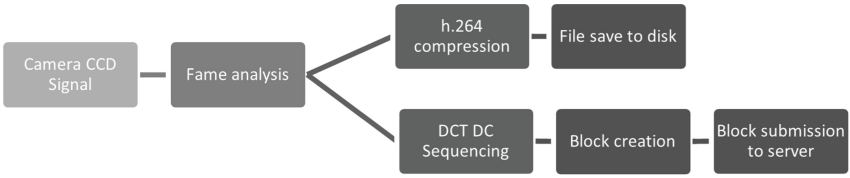

Fig. 1. Multi-threaded processing of frame for blockchain and video pipeline.

All of these options have some applicability and could be utilised to some effect. We have prioritised location based integrity checking with a minimum impact on image quality, and therefore combined a QR code with a DCS operation on keyframes as our method of identification and integrity checking. Our blockchain implementation utilises a Merkle tree structure that consists of blocks created on the camera in real time. Our previous work led to the adoption of a tree structured to create independent hash branches for each camera, merging into the root upon submission to the server, this flexibility is useful in some LEA operations where the camera must operate independently for some time

before retrieval. Figure 2 shows the Merkle tree of two separate cameras generating their own blockchain records, with each block identifying hash, H_i, being hashed with the previous root hash, R_{i-1}, to form the Merkle root (R_i) of the system.

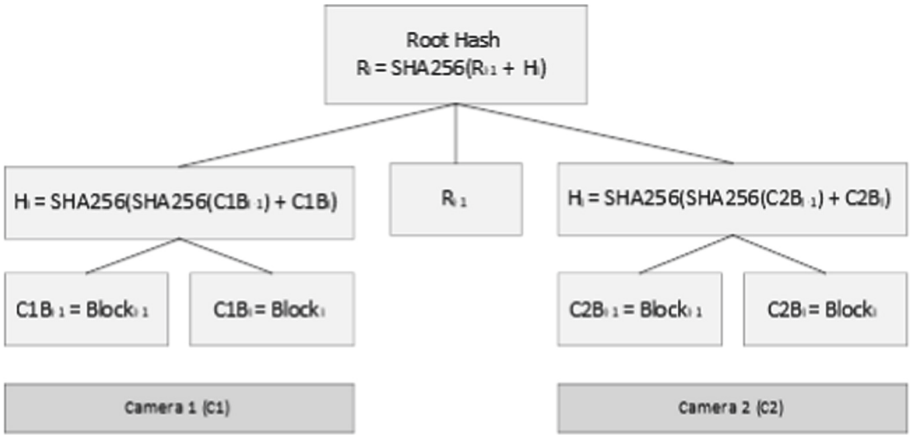

Fig. 2. Hash tree structure allowing for block records to be submitted in real time from multiple cameras, C1 and C2.

Figure 3 provides an overview of an operational system. Video is collected in the field and blockchain records are immediately transmitted to the DLT primary server using either the cameras own radio, or a tethered cellular or LMR communications device. The server infrastructure replicates DLT records, allowing for access from outside LEA networks. Third party clients can access and query the block's metadata using any DLT server without having access to the archived video.

1.3 Reference System

Our reference system contains the main components of Fig. 3. Our rugged wearable recording device utilises Wi-Fi to connect to a mobile bearer connecting to the Primary DLT server. Files are recorded locally and uploaded to central storage representing an end-of-shift LEA archive process. The Ledger is replicated to a Secondary DLT server, representing a 3rd party oversight authority. Our bespoke client application is used to search for blockchain records and verify both video and blockchain integrity. We examine the effectiveness in four distinct scenarios.

Scenario 1 (S1). Normal operation. Officer records video in the field. Each block creation event produces a block that is transmitted to the DLT primary server in real time.

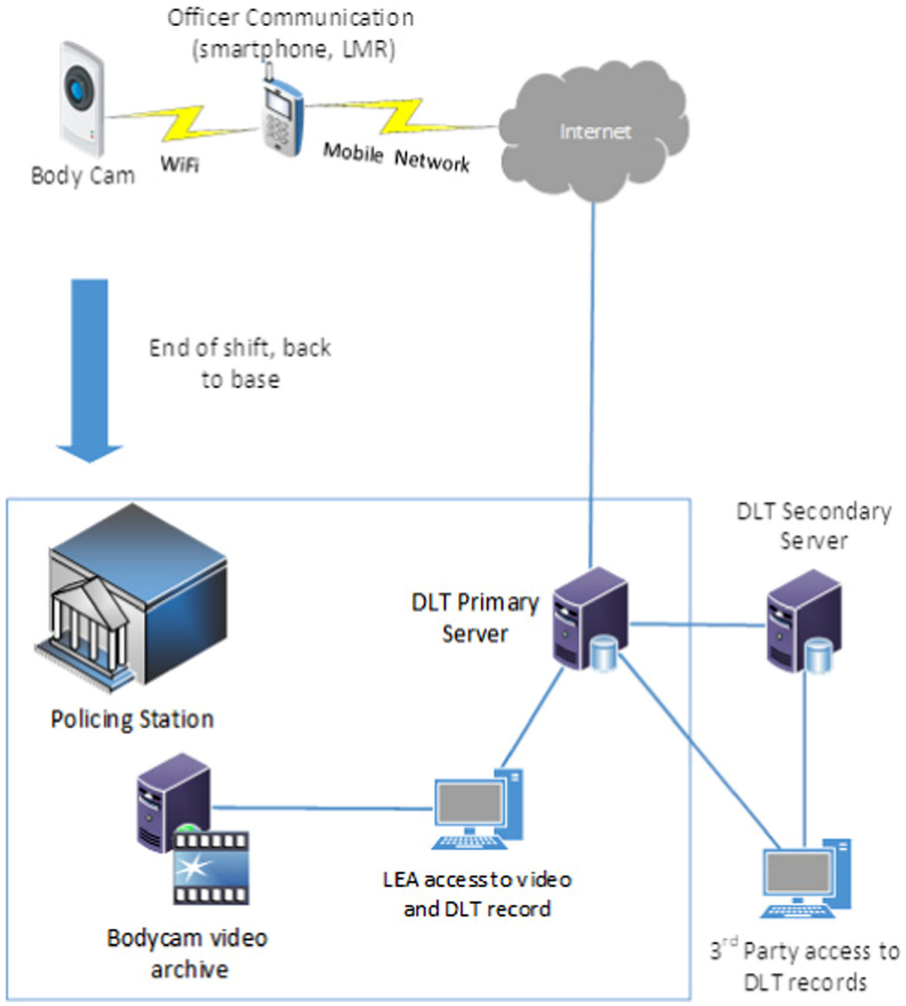

Fig. 3. System architecture overview.

As shown in Fig. 2, for Camera 1 (C1), the block hash $C1H_i$ is hashed with the Root, R_{i-1}, to produce R_i. We conducted three recordings of manually triggered events (E1, E2 & E3). A basic collection of metadata was added to each block: CameraID (C), BlockID (B), Creation Time (CT), DC Sequence (DCS), Event Code (E) the Block hash (H_i), being Eq. 1.

$$H_i = SHA256(C + B + CT + DCS + E + H_{i-1}) \tag{1}$$

Our Bodycam hardware consisted of a generic embedded device running Linux, capturing video at 15 fps, 640×480 continuous for 5 min. The format and length of S1 sessions is summarised in Table 1.

Table 1. Collected data from S1.

Event	Block count	S1 Block Avg kb	DCS Avg kb	Duration
E1	319	28.898	28.567	5 m:22 s
E2	324	28.568	28.237	5 m:28 s
E3	326	28.572	28.241	5 m:30 s

Scenario 1a (S1a). Normal Operation. The workflow is identical to S1, except the DCS is not submitted in real time. This data is backfilled once the camera has returned to base. The output can be derived on the data collected in S1 as *S1 Block Avg kb - DCS Avg kb*.

Scenario 2 (S2). Normal Operation. Given a video clip, locate and identify its blockchain data and verify the video integrity.

Given a known date and time, event code or operational identifier, blockchain queries are trivial, they are not covered here. For S2, searches were conducted in the reverse direction, given a clip of an arbitrary length, the blockchain is queried for the DCS data, using Pearson's correlation as our metric, our bespoke client tool produced both tabular and visual feedback on the clip's correlation to the blockchain stored DCS, grouped by keyframe. Through the use of client side object detection [11] granularity was further improved by grouping DCS correlation by object. We then calculated the Mean Deviation to provide an indicator of consistency of the comparison throughout the clip by object, these results are summarised in Table 2.

Table 2. Results of S2 client processing.

Event	Frame correlation mean	Frame mean deviation	Object correlation mean	Object mean deviation
E1	0.9989	0.0004	0.9996	0.0003
E2	0.9988	0.0005	0.9983	0.0013
E3	0.9994	0.00008	0.9977	0.0020

Scenario 2a (S2a). Abnormal Event. Perform the same procedure as S2, but with an altered video clip.

E1 video collected in S1 is clipped and colour tinted. When reprocessed the correlation score is much lower relative to known good samples. For example, the Mean Deviation of our example frame in Table 3 is .0105 (.9883 − .9778), higher than standard deviation for object level correlation of the sample, and much higher than the averages recorded in Table 2. A sample original and altered frame is shown for reference in Table 4.

Table 3. Analysis of the modified sample clip produced by S2a.

Event	Frame correlation	Frame mean deviation	Object correlation mean	Object mean deviation
E1	0.9913	0.0083	0.9883	0.0086

Table 4. Frame examples from an original and a modified video clip.

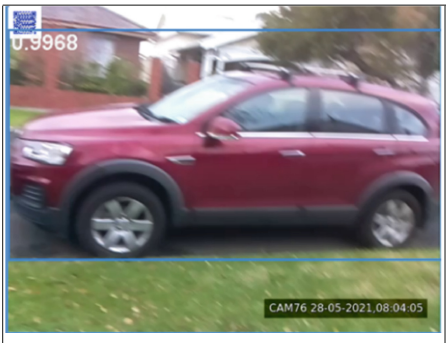	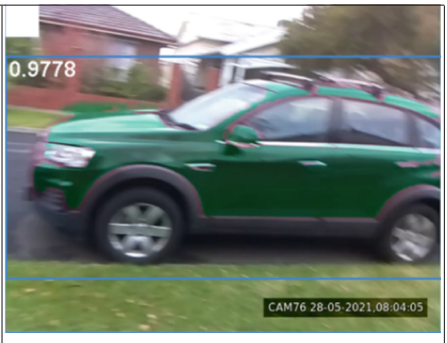
Original Video correlation for 'car' object .996	Modified Video correlation for 'car' object .9778

1.4 Results Discussion Points

Blockchain Transmission Efficiency. Table 1 provides a comparison between the DCS data field size, and the entire blockchain message including all described metadata fields and the AMQP (Advanced Message Queue Protocol) overhead, it is clear the S1a strategy provides an improvement of approximately 26 kilobytes per block (when transmitting uncompressed AMQP messages) by real time transmission of partial data. This comes at the expense of data being unavailable in the blockchain until it is uploaded (i.e., at the end of the officer's shift).

Video Identification and Integrity Analysis. Averages for the correlation of all three video segments were very high, (min .998), additionally the Average Deviation from Mean remained very low, not higher than .002, indicating consistent operation and verification of the video clip.

Analysing the correlation Mean of objects within the video and the Mean Deviation allows for more granular analysis. Recorded values also followed the pattern of the frame analysis (Table 2). It is important to note in this implementation of YOLO [11] objects detected in subsequent frames are not identified as the same real world object, therefore single objects appear and are evaluated as many times as they are detected within the clip.

For S2a, the resulting values in Table 3 show that modified objects can be detected by overall lower correlation average scores as well as higher than normal Standard Deviation of Mean values, this value however would be dependent on factors such as the length and number of objects detected, and the number of modified objects and the length of the sample clip.

1.5 Further Work

Of interest is the reduction of the DCS field size. Implementing zLib compression yields between 2:1–5:1 compression [12], and further to this additional strategies such as identification and storage of deltas only are currently being developed. When utilising the S1a partial upload strategy there is a risk that if the Bodycam fails to eventually upload the entire block, H_i cannot be recreated and the chain can no longer be verified, either against itself or against the Merkle root. Operational circumstances may create such a situation if the camera is damaged or lost in the field.

In order to implement the benefits shown in S1a the Merkle tree structure should be extended to support a two stage hash algorithm, allowing for recovery from the partial upload of blocks and the subsequent loss of destruction of the camera. Whilst not ideal this will at least provide an audit record, and so regardless of the workflow surrounding the management of video clips, the collection of recordings are immediately a public record.

Bodycams are almost always recording audio along with the video stream, currently there is no provision for audio data integrity protection. In much the same way as we have viewed watermarks as interchangeable, watermarking of audio streams would be similarly approached, linked to the same DLT infrastructure.

Where possible the system utilises infrastructure security features, such as transport SSL encryption provided by the Message Queue software. Whilst write access to the blockchain is currently controlled by secret key this does not address the potential for repudiation of blocks after they are committed. Therefore, subsequent versions of the camera software will support digital signing of the blocks on the camera.

Our system performs best when it can control the creation of keyframe data; although every effort has been made to implement a cross platform, hardware agnostic on-camera process, we have found inconsistency in some multimedia

implementations on how keyframes are triggered and specified. Further work is required here, and it may be unavoidable that some customisation may be necessary to accommodate some system APIs.

1.6 Conclusion

Our software is designed to process large amounts of CCTV video data and produce statistical indications of integrity and abnormal operation. It is demonstrated here to another LEA requirement of cataloguing and verification of body worn video data. The Standard Deviation and Mean Deviation functions provide a method to process bulk data to detect anomalies of tampering, missing data or device malfunction. This method may not always be precise enough to conclusively determine tampering, but it is valuable as a triage method, and a method to recommend further analysis, especially for 3rd parties without access to original recordings, such as civil liberty groups. The system shown here is lightweight and vendor agnostic, allowing for implementation on a range of low powered collection devices and varying police workflows. Such a system would contribute to growing trust in police body worn surveillance systems and recognition of their potential to protect the interests of both civilians and policing officers alike.

References

1. Doyle, A., Lippert, R., Lyon, D.: Eyes Everywhere: The Global Growth of Camera Surveillance. Taylor & Francis Group, London (2012)
2. Goold, B.J.: Not just about privacy: police body-worn cameras and the costs of public area surveillance. Police Camera: Surveill. Priv. Accountability **2020**, 167–181 (2020)
3. Marx, G.T.: Introduction: The ayes have it - Should they? Police Body-worn Cameras (2021)
4. Blanchette, J.-F., Becker, S.: Bodycam footage as document: an exploratory analysis. In: Chowdhury, G., McLeod, J., Gillet, V., Willett, P. (eds.) iConference 2018. LNCS, vol. 10766, pp. 609–614. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78105-1_68
5. Police Body Cameras—American Civil Liberties Union. https://www.aclu.org/issues/privacy-technology/surveillance-technologies/police-body-cameras. Accessed 07 May 2021
6. EIP-20: ERC-20 Token Standard. https://eips.ethereum.org/EIPS/eip-20. Accessed 21 May 2021
7. Kerr, M., van Schyndel, R.: Adapting law enforcement frameworks to address the ethical problems of CCTV product propagation. IEEE Secur. Privacy **12**(4), 14–21 (2014)
8. Kerr, M., Han, F., Schyndel, R.V.: A blockchain implementation for the cataloguing of CCTV video evidence. In: Proceedings of AVSS 2018–2018 15th IEEE International Conference on Advanced Video and Signal-Based Surveillance (2019)
9. Chandrakar, N., Bagga, J.: Performance comparison of digital image watermarking techniques: a survey. Int. J. Comput. Appl. Technol. Res. **2**(2), 126–130 (2013)

10. Gstreamer. Gstreamer Project (2021). http://gstreamer.freedesktop.org. Accessed 24 May 2021
11. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: unified, real-time object detection. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, vol. 2016-December, pp. 779–788 (2016)
12. Zlib Technical Details. https://zlib.net/zlib_tech.html. Accessed 24 May 2021