Tamper-Proofing Imagery from Distributed Sensors Using Learned Blockchain Consensus

Alexander D. Wissner-Gross
Gemedy, Inc. &
Institute for Applied Computational Science
Harvard University
Cambridge, MA, USA
Email: alexwg@post.harvard.edu

Jared C. Willard
CCDC Army Research Laboratory
Aberdeen Proving Ground, MD, USA
Email: jwillard211@gmail.com

Noah Weston
CCDC Army Research Laboratory
Aberdeen Proving Ground, MD, USA
Email: noah.d.weston.civ@mail.mil

2020 IEEE Applied Imagery Pattern Recognition Workshop (AIPR) | 978-1-7281-8243-8/20/$31.00 ©2020 IEEE | DOI: 10.1109/AIPR50011.2020.9425050

***Abstract*—Area monitoring using wireless sensor networks that collect imagery and multimodal data from multiple vantage points, while requiring only limited local bandwidth and compute resources, promises improved resilience and scalability over single-camera imagery. However, the distributed nature of such networks can also increase their relative vulnerability to subversion via physical tampering. Here we address that nascent vulnerability by introducing a blockchain application that (1) learns correlations between low-dimensional projections of observed sequences captured by pairs of sensors, and (2) uses those correlations as a baseline for a soft consensus mechanism that identifies potentially compromised sensors with the strongest pairwise statistical anomalies. We then demonstrate our approach in a simulated environment in which a network of virtual cameras are aimed at a common dynamical scene from different vantage points, and show that after a training period of observing baseline behavior followed by the subversion of various numbers of the cameras, our application can correctly identify the cameras that have been compromised. Finally, we explore automated responses to such compromised sensors, including denying them shared resources and services.**

1. Introduction

Successful operations in future contested environments are expected to rely heavily upon autonomous agents collaborating among themselves and with humans [1]. In particular, wireless sensor networks that collect imagery and multimodal data from multiple vantage points, while requiring only limited local bandwidth and compute resources, represent a class of such agents with promising resilience and scalability characteristics. However, the distributed nature of such networks can also increase their relative vulnerability to subversion via physical tampering.

How can we enable such decentralized area monitoring systems to identify, and even tolerate, anomalous behavior from potentially subverted camera sensors? We may draw inspiration from the more prosaic problem of Byzantine Fault Tolerance (BFT), which was solved by Shostak, Pease, and Lamport [2] with algorithms that are now widely adopted in blockchain systems and real-time aircraft systems, among other applications. BFT requires $3N + 1$ nodes to achieve correct consensus in the presence of N malicious nodes, or $3N$ nodes with digital signatures. For finely grained, unsigned-sensor anomaly detection and correction, that would mean $3 \times 1 + 1 = 4$X redundant image coverage of monitored areas. If we arranged cameras naively in collocated 4-clusters, as shown in Figure 1(a), we could achieve local BFT, but each sensor would be vulnerable to common physical attacks on its cluster, negating the benefits. To address that issue, a better arrangement for 4X coverage would be a mesh of cameras with overlapping fields of view, as shown in Figure 1(b). However, if camera groups share overlapping coverage areas, but do not generate replicated data due to their differing vantage points, as illustrated in Figure 1(c), can fault tolerance still be achieved? Here we report a solution to that problem by implementing a blockchain smart contract that learns correlations between camera streams and uses those correlations as a baseline for a soft consensus that identifies cameras with the most pairwise statistical anomalies, as illustrated in Figure 1(d).

2. Theory

Our theoretical approach to achieving soft consensus among imagery sensors assumes that we capture at least one frame of video per camera i per blockchain block t, and compute a scalar projection $c_{i,t}$ of the frame to avoid needing to store entire images on-chain. For simplicity, we chose $c_{i,t}$ to be the sum of pixel intensities over RGB channels of each image.

To accommodate differing vantage points, we focused on tracking pairwise linear correlations between camera image projections, since dynamic environments will often present visual anomalies at the single-camera level, but historical correlations between cameras with overlapping fields of view should be stationary under normal conditions. For simplicity, we further assumed camera positions whose image projections would be positively correlated, allowing us to focus on detecting anomalous differences $d_{i,j,t} = c_{j,t} - c_{i,t}$

978-1-7281-8243-8/20/$31.00 ©2020 IEEE

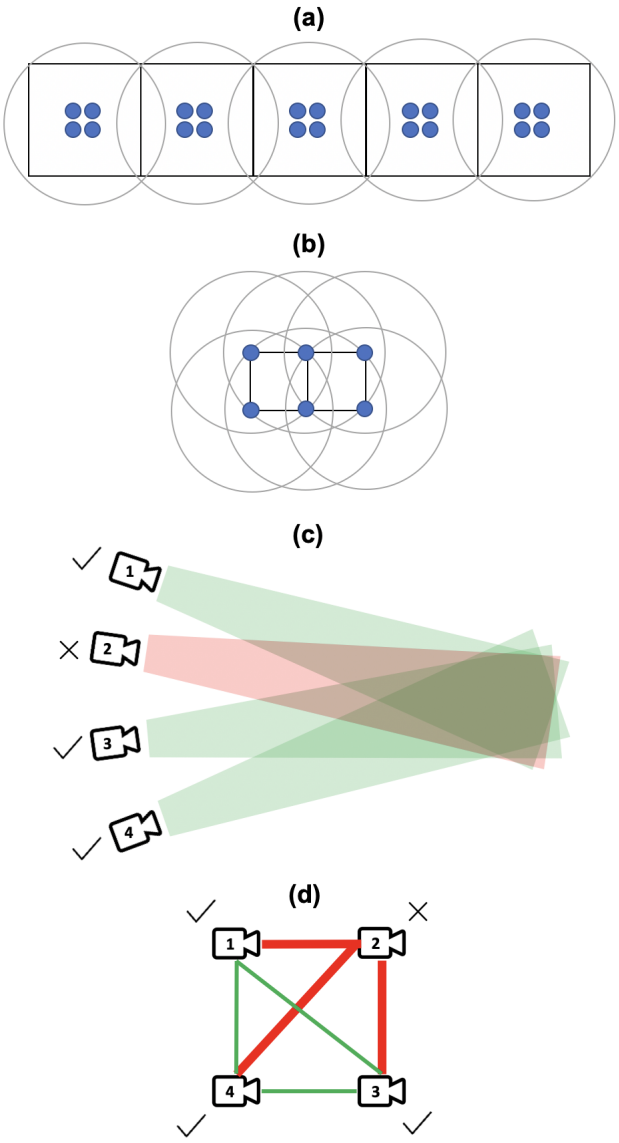

Figure 1. Schematic illustrations motivating our approach to fault-tolerant area monitoring. (a) Naive arrangement of cameras into collocated 4-clusters. (b) Mesh of cameras with 4X overlapping coverage. (c) Soft consensus among cameras with different vantage points and one subverted camera. (d) Our approach to detecting subversion through correlation statistics between pairs of cameras.

between cameras i, j. A more general approach for the future might utilize Mahalanobis distances [3] to allow for negative correlations as well.

When a new projection from camera i was reported, we calculated new pairwise z-scores between cameras $i, j \neq i$:

$$z_{i,j,t} = \frac{d_{i,j,t} - \frac{1}{t}\sum_{t'<t} d_{i,j,t'}}{\sqrt{\frac{1}{t}\sum_{t'<t} d_{i,j,t'}^2 - \frac{1}{t^2}\left(\sum_{t'<t} d_{i,j,t'}\right)^2}}$$

TABLE 1. SIMULATED SCENARIOS USED FOR VALIDATION

Scenario	Name	Description
I	Normal	A vehicle moves through an intersection and is observed by all four cameras.
II	False Negative	A vehicle moves through an intersection and is observed by three cameras, but the fourth camera has been subverted and sees an empty road.
III	False Positive	An intersection is empty and is observed as such by three cameras, but the fourth camera has been subverted and sees a vehicle.
IV	Partial Occlusion	A vehicle moves through an intersection and is fully observed by three cameras, but the fourth camera is partially blocked and cannot see the full vehicle.

using only blockchain-compatible integer arithmetic:

$$z_{i,j,t}^2 = \frac{\left(d_{i,j,t}t - \sum_{t'<t} d_{i,j,t'}\right)^2}{t\sum_{t'<t} d_{i,j,t'}^2 - \left(\sum_{t'<t} d_{i,j,t'}\right)^2}.$$

By convention, we considered pairwise z-scores above the threshold of 3-sigma to be anomalous. For $N = 4$ overlapping sensors, a total of at least $2(N-1) = 6$ pairwise anomalies over two consecutive blocks involving a given camera i could therefore be interpreted as a soft consensus that camera i's individual behavior was anomalous.

3. Experiment

To validate our approach to learned soft consensus for fault-tolerant distributed area monitoring, we used the ANVEL simulation environment [4] to explore four experimental scenarios, as enumerated in Table 1. In the first scenario (I), a vehicle was manually driven through a multiway urban intersection overseen by a vertical tower of four normal cameras with overlapping but distinct fields of view, as shown in Figure 2. In the remaining three scenarios, a single camera was subverted to falsely report no vehicle (II), one vehicle (III), or a partial vehicle (IV), respectively. Camera oracles reported projections of their 640×480 image captures every ca. 5 seconds, matching the 5-second block time.

For the supporting blockchain, we used ARL's Tactical Distributed Ledger (TDL), an extensible tiered computing framework of decentralized Ethereum [5] applications overlaid on a peer-to-peer network. TDL allows participants to register available services and capabilities, request services of each other, dynamically deploy new capabilities as the need arises, and collaborate on computing tasks in an ad-hoc environment. The anomaly-detection contract represented just one module within a web of smart contracts forming the TDL, as shown in Fig 3.

Figure 2. Vertical tower of angled cameras (seen on the left) overlooking a vehicle in a multiway urban intersection, as simulated in ANVEL.

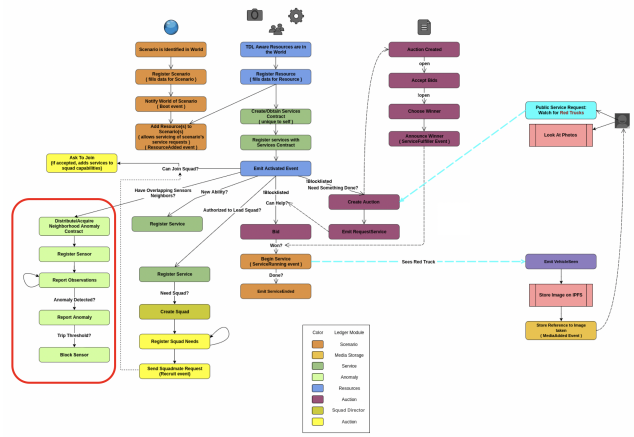

Figure 3. Schematic illustration of the ARL Tactical Distributed Ledger (TDL) web of smart contracts, including the anomaly-detection contract reported here (circled in red).

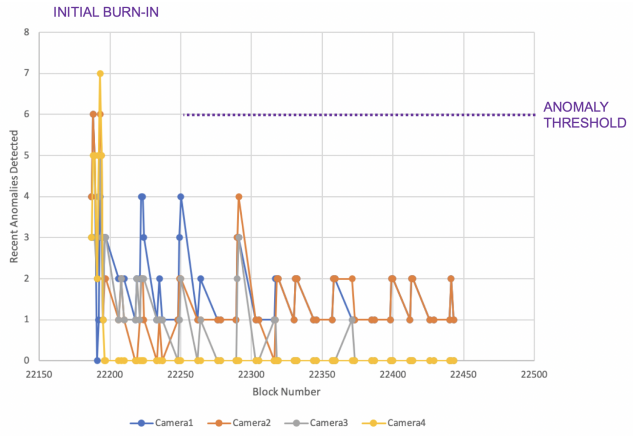

Figure 4. Recent anomalies detected over time (block) for Scenario I (Normal).

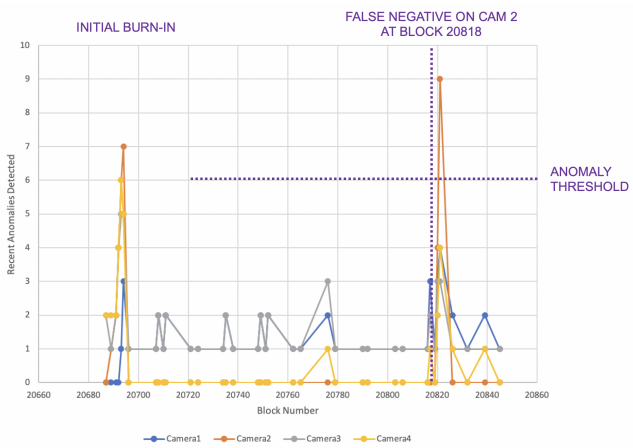

Figure 5. Recent anomalies detected over time (block) for Scenario II (False Negative).

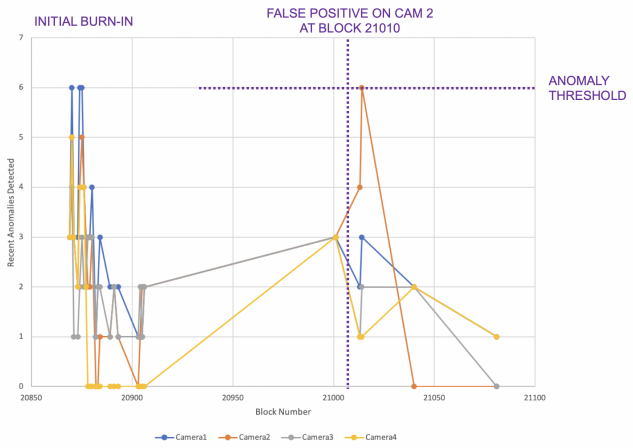

Figure 6. Recent anomalies detected over time (block) for Scenario III (False Positive).

4. Results

In every scenario that we considered, after an initial "burn-in" training period of several blocks, during which normal pairwise camera correlations were learned, the smart contract was able to correctly identify the absence (Scenario I) or presence (Scenarios II, III, and IV) of subversion. Moreover, in the presence of subversion, the smart contract correctly identified the specific camera that was reporting manipulated video. In the normal Scenario I, following initial burn-in, the theoretical threshold of 6 recent pairwise anomalies was never reached (as seen in Figure 4). In contrast, that threshold was reached immediately upon subversion of a single camera by, and only by, recent pairwise anomalies for that camera in Scenarios II (Figure 5), III (Figure 6), and IV (Figure 7), demonstrating the sensitivity and specificity of our approach.

5. Conclusions

We have presented a new approach to tamper-proofing of imagery, in which statistical learning was combined with a blockchain-based consensus mechanism to detect pairwise anomalies between manipulated and unmanipulated camera

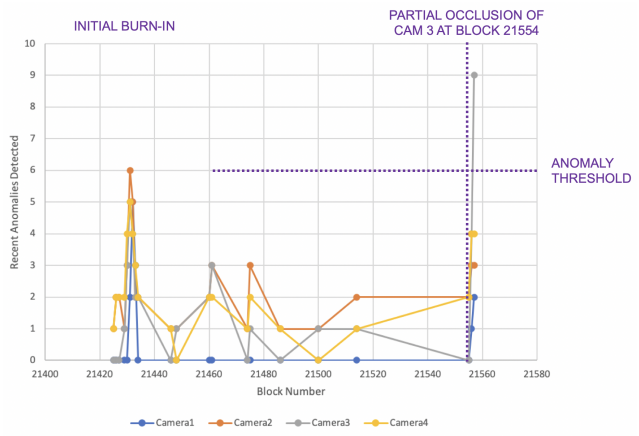

Figure 7. Recent anomalies detected over time (block) for Scenario IV (Partial Occlusion).

streams. We demonstrated our approach in a simulated physical environment in which networked virtual cameras were aimed at a common dynamical scene from different vantage points. We showed that after a training period of observing baseline behavior followed by the subversion of a camera, our system could correctly identify the camera that had been compromised. Our approach opens the door to automated responses to subverted image sensors, including denying them shared resources and services. More generally, our approach to tamper-proofing imagery using learned blockchain consensus promises to improve the resilience and scalability of area monitoring.

References

[1] A. Kott, "Challenges and characteristics of intelligent autonomy for internet of battle things in highly adversarial environments," in *AAAI Spring Symposium Series*, 2018, pp. 145–151.

[2] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," *ACM Transactions on Programming Languages and Systems*, vol. 4, no. 3, pp. 382–401, 1982.

[3] P. C. Mahalanobis, "On the generalized distance in statistics," *Proceedings of the National Institute of Sciences of India*, vol. 2, no. 1, pp. 49–55, 1936.

[4] P. J. Durst, C. T. Goodin, C. L. Summins, B. Q. Gates, G. B. McKinley, and T. R. George, "A real-time, interactive simulation environment for unmanned ground vehicles: the autonomous navigation virtual environment laboratory (anvel)," in *Fifth International Conference on Information and Computing Science*, July 2012, pp. 7–10.

[5] G. Wood, "Ethereum: A secure decentralised generalised transaction ledger," *White Paper*, 2014.