A Blockchain-Based Data-Sharing Scheme for Inter-vehicular Safety Applications

Doug Lundquist(✉)

University of Illinois at Chicago, Chicago, USA
dlundq1@uic.edu

Abstract. Vehicular safety applications could save lives by sharing data not available from line-of-sight sensors but they also require trust among a set of mutually distrustful vehicles. We present a scheme for sharing validated vehicular trajectory data via vehicle-to-vehicle communication to help reduce traffic collisions. It does not require any centralized control or roadside infrastructure to function. Instead, vehicles share and validate data directly among each other. Our scheme combines a distributed blockchain model to create a permanent set of validated trajectory data. Vehicles join one or more consortium blockchains shared among nearby vehicles. Within each blockchain, vehicles share data between others nearby through a fully decentralized controlled flooding protocol. As blockchain and vehicular networks are prone to scalability concerns, we have designed our scheme specifically to address them. It limits the number of vehicles participating in each blockchain, bounds how widely trajectory data are shared, and organizes and merges redundant data to reduce total network traffic. We also discuss several future directions for assessing the relative performance profiles of specific blockchain and networking implementations.

Keywords: V2V applications · Blockchain · Vehicular safety

1 Introduction

Widespread adoption of inter-vehicular safety applications (IVSAs) could save thousands of lives lost to traffic collisions every year [1] using vehicle-mounted mobile devices to share information among nearby vehicles. Although vehicles can use their own line-of-sight sensors to detect and evaluate dangers, many vehicular collisions occur specifically because the colliding vehicles cannot see each other. A common example is one car with the right-of-way entering an intersection while another, concealed behind a large truck, runs a red light and crosses its path (Fig. 1). Such collisions could be prevented by vehicles sharing their real-time trajectories by wireless communication with others nearby.

For IVSAs, accurate and secure information is vital. Malfunctioning sensors could create and distribute bad data. Likewise, malicious participants might deliberately alter IVSA data to provoke vehicular collisions. Many mechanisms to track reputation and

W. Bao et al. (Eds.): ADHOCNETS 2021/TridentCom 2021, LNICST 428, pp. 205–213, 2022.
https://doi.org/10.1007/978-3-030-98005-4_15

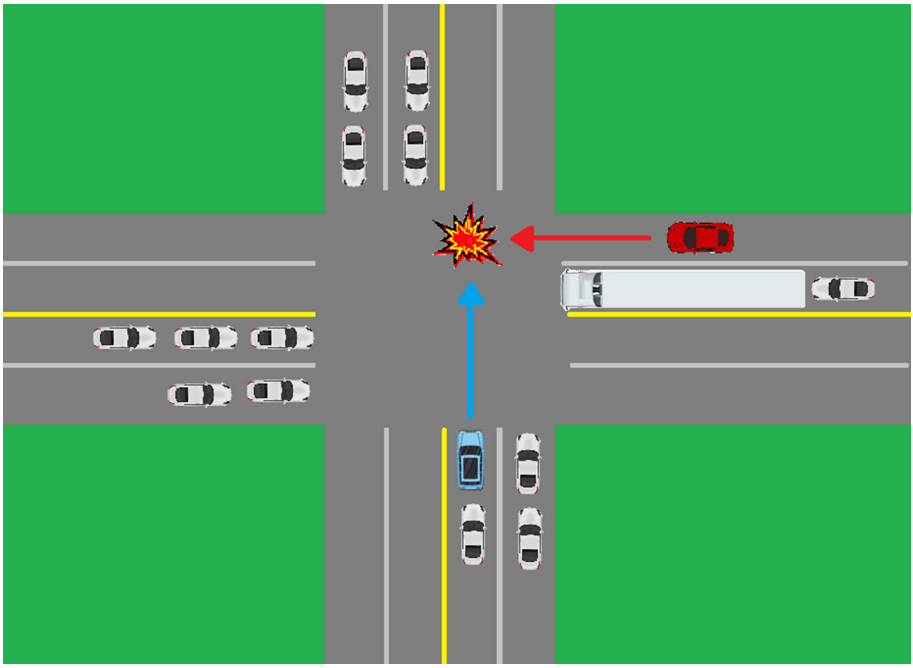

Fig. 1. Traffic collision at an intersection: the blue vehicle has the right-of-way but does not see the red vehicle entering the intersection. (Color figure online)

discourage good behavior have been proposed - [2] provides a good survey - but the recent development of blockchain technology offers novel solutions to the challenges of trust. In blockchain databases, batches of new data are confirmed by mutual agreement. Rather than guaranteeing good behavior, blockchain creates trust between participants because any improperly updated data would be easily detected. This offers a strong defense against isolated malicious participants - vehicles will get the correct information from other, non-malicious sources.

This paper proposes a peer-to-peer data management framework for blockchain-based IVSAs without infrastructure or a stable set of vehicles known to be trustworthy. All secure blockchain models require that either certain users are trusted or the blockchain is shared by a large user group to prevent its takeover by malicious actors. In IVSAs without infrastructure, however, participation of trusted vehicles would be difficult to guarantee. Thus, we focus on blockchain models with distributed ownership. A distributed blockchain - shared by users with equal standing - requires a fairly numerous and continuous population of users to provide data security.

Although the raw number of participants is readily obtained in IVSAs, group membership must still be managed. For example, a group of nearby vehicles belonging to a blockchain would periodically lose members over time due to divergent or completed travel paths. In fact, it is entirely possible that an IVSA could use a hybrid blockchain model. For example, it could combine private and public elements or proof-of-stake intermittently supported by proof-of-work. Likewise, vehicles could simultaneously belong

to multiple groups to help ensure a continuous data flow. In this paper, we are therefore agnostic about the specific blockchain implementation and how it manages group membership.

The core of our framework is managing data flows within the vehicular network. Our goal is to ensure fast delivery of useful application data, e.g., the trajectories of nearby vehicles. In large-scale IVSAs, delivering every vehicle's data to every other vehicle is undesirable, as the heavy network traffic would delay delivery of immediately useful data. In large-scale systems, propagation limits must either be enforced or else arise organically and unpredictably when network traffic injections exceed the actual delivery capacity. We propose adopting the Self-Balancing Supply/Demand (SBSD) protocols [3], which dynamically bound data propagation according to metrics of the data's age, popularity, and distance from the source.

The rest of this paper is structured as follows. Section 2 discusses a selection of relevant research. Section 3 briefly covers the basics of blockchain and the SBSD model. Section 4 describes how our scheme will operate, in particular the mechanics of generating, processing, and sharing data. We conclude in Sect. 5 with a summary of our proposal and offer a few future directions for our research.

2 Related Research

Useful vehicle-to-vehicle (V2V) communication technologies exist but have not been widely implemented. Basic questions, like the selection and usage of wireless communication standards and security policies, have not been fully unresolved. Ongoing debates involving automakers, technology developers, and government agencies – essentially, what technology to use and how to use it - have stalled the adoption of existing technology. In fact, in the United States, half the bandwidth reserved for V2V communications was recently released for general usage [4].

The recent, explosive growth of the cryptocurrency sector has followed decades of research in data security that eventually led to blockchain. Interested readers might consult [5] to learn about core blockchain models and their applicability to real-world problems. Certainly, blockchain has received a great deal of attention for processing financial transactions, reflected in an abundance of research output. Security is naturally also a core concern for blockchain research; [6] provides a good introduction to the major security issues in blockchain implementations. In vehicular contexts, researchers have proposed and evaluated adopting blockchain for various building blocks of secure IVSAs, such as:

- Data management via blockchain in vehicular ad hoc networks [7, 8]
- Securely processing payments between electric vehicles and the grid [9]
- Privacy-preserving authentication [10–12]
- Group management: managing blockchain group memberships [13, 14]
- Evaluating the performance of blockchain under the high mobility typical of vehicular networks [15].

The other element of our scheme is the SBSD model for sharing data. SBSD is an alternative to models that manage routing paths in vehicle groups. By including age

and distance metrics with packets, SBSD regulates data propagation without directly managing vehicle groups. Compared to [14], no clusters need to be managed and no cluster heads elected. Similarly, in [13], the blockchains are linked to large spatial areas without a mechanism for dynamically changing the covered areas. In contrast, our SBSD model is inherently adaptable to changing conditions of vehicle population, effective transmission distance, and wireless data transmission capacity.

SBSD grew out of an earlier model [16] for ranking data according to factors like age and popularity, as a mechanism for delivering the most interesting data within transmission capacity limits. When data flows exceed network capacity, participants must decide which packets to forward. In SBSD, only high-ranking packets (*e.g.*, new and close their point of origin) are forwarded by receiving vehicles. Analogous to Facebook's EdgeRank algorithm, which ranked the order in which posts would appear in a user's feed, SBSD uses similar methods to regulate data propagation.

3 Model Components

Let us now briefly cover the basic concepts of blockchain and SBSD, in Subsects. 3.1 and 3.2, respectively.

3.1 Blockchain

Fundamentally, blockchain is a data management technology for letting mutually distrustful users share information and quickly verify its correctness. Data is managed in batches called blocks. When a new block is created, database records can be added to it by blockchain users. However, these records must be approved by consensus - if one user tries to insert bad data, other users are expected to reject it and prevent its insertion. In IVSAs, vehicles would advertise their own recent trajectory and others would confirm its correctness. Eventually, the current block will be full - they hold a finite number of records - and it will be permanently added to the blockchain, with a timestamp and hash codes to facilitate detection of any changes to the finished block.

Next, we consider who can use a given blockchain. Blockchains can be public – allowing anyone to read and write data – or private, only allowing access to members of a single organization. An intermediate possibility is the consortium blockchain model, which allows access to members of multiple organizations. However, the larger the pool of users, the more work is entailed in getting a majority of users to approve a record. To ensure the processing needed to prevent vehicular collisions can happen in real-time, each blockchain would be owned by a small, localized vehicle group.

Within the IVSA context, each new blockchain database record would be a fixed-time trajectory segment and each block would let every vehicle add the same number of records. Every block would thus cover a known time interval and every record in a parent block would predate every record in any of its child blocks. This would facilitate recognizing missing records in the block, helping ensure a complete and sequential set of traffic data.

To create a new record, a vehicle announces a set of its own recent trajectory segments. When other nearby vehicles confirm the set, confirmed record would be shared with the

rest of the group and be added to their block copies. Within an appropriate time after the expected end of the block's time interval, the block would be confirmed by consensus. This blockchain model would give vehicles actionable data before finalizing each block. In the long term, sources of improperly altered data can be recognized. In cases of disputed data, vehicles might cautiously treat all conflicting data as being possible, but this policy certainly requires further development.

Blockchain does present unique security risks since consensus can be gained through fraud. For example, a single vehicle controlling many phony vehicle identities could confirm bad data. Likewise, a set of malicious users with a simple majority of the group could achieve the same outcome, i.e., a 51% percent attack. Although it would be possible for a trusted authority (such as roadside infrastructure) to provide this security, we envision a system distributed over the vehicles themselves. In this case, a trusted blockchain should have its ownership shared over a large group of users over time, with reliable user verification.

Because blockchain is a novel technology with standards still in development, there are many questions about how a blockchain system with adequate security and speed would be implemented. We do not address that question in this paper nor claim superiority for a particular blockchain implementation, *e.g.*, in the proof-of-work vs. proof-of-stake approaches. In fact, a single system might well use multiple blockchain technologies.

For example, hybrid systems like Decred employ both proof-of-stake (for its low energy consumption) and proof-of-work (to offer security against large but malicious cryptocurrency holders). Accordingly, we believe the membership application should be a modular one, which simply gathers vehicle population data and, when needed, forwards that information to the various blockchain users. Ideally, this application should also be able to automatically regulate bandwidth consumption among a background of other V2V applications.

3.2 Self-balancing Supply/Demand

The SBSD framework provides a low-overhead probabilistic framework for regulating network traffic flows. Fundamentally, SBSD is a controlled flooding protocol which dynamically limits every packet's flooding extent. To limit propagation, vehicles only forward (via broadcast) the most relevant packets within the limits of their transmission capacity. Each packet's header stores its age and hop count from its source to allow calculating the relevance metric:

$$1/\left[(age)(hops)^{1}/_{2}\right] \tag{1}$$

As a packet travels away from its source, each hop increases the hop count and the packet gets older, decreasing the packet's relevance and making it less likely to be forwarded again.

The relevance metric may be multiplied by a popularity factor, frequency, which tracks the number of times a packet is independently created. For example, an urgent safety message might be simultaneously created by multiple vehicles. Frequency will multiply the relevance metric above, *i.e.*:

$$(frequency)/\left[(age)(hops)^{1}/_{2}\right] \tag{2}$$

More popular content has higher relevance, all else equal, causing those packets to be forwarded for longer and over a larger area. For a blockchain group, frequency could also be aligned with group populations and areas, so that packets are flooded over the approximate extent of the group.

Finally, to hasten packet delivery in the network, new packets are transmitted according to a binary exponential backoff model. A packet received at a node n with enough relevance to be forwarded will typically be forwarded at n's next opportunity to transmit. Thereafter, subsequent retransmissions will take about twice as long. This approach helps ensure that new data injected into the network will be quickly shared but retransmissions will substantially slow down after most nodes in the vicinity have probably already received it.

4 Model Operation

This section describes how our scheme shares and processes vehicle trajectory data in a blockchain. We assume that the blockchain mechanics – like tracking membership and finalizing each block – are handled outside of our scheme. We also note that our scheme is inherently designed to accommodate vehicles simultaneously belonging to more than one blockchain.

In our scheme, vehicles maintain a table of data for each blockchain, consisting of a list of vehicles, their trajectories, and the sources of that data. We describe the data vehicles generate and transmit (Sect. 4.1), then how recipients process the data (Sect. 4.2), how the data may be forwarded (Sect. 4.3), and beneficial aspects of our approach (Sect. 4.4).

4.1 Generation

Each vehicle will periodically measure and transmit its own trajectory to other nearby vehicles, with itself as the source. Let the transmitting vehicle be v_s and the set of recipients be V. The source v_s will add its own trajectory t_s to the tables for its blockchains. The initial relevance of the packet will be created with *age* and *hops* at *0*. Because the packet will have high relevance, it will be transmitted immediately and then given priority in forwarding.

4.2 Processing

When a vehicle v_r in V receives the packet from v_s, it will compare the new data to its own last confirmed trajectory from v_s. If the data appears correct, it will add its own identifier to the packet payload. Ideally, the data verification would use line-of-sight but alternative methods could be used for vehicles that cannot be seen due to terrain obstacles. For example, v_r could verify that the recent trajectory segments connect and any speed changes are plausible for v_s.

The updated trajectory data would be kept as a tree list of verifications for the trajectory. For each vehicle in the blockchain, v_r would maintain entries as a list starting with the root – the source vehicle v_s – and the sequence in which verifications were

made. In the tree graph, all vehicles in V would be child nodes of v_s and each vehicle in V could have its own child nodes, vehicles which did not receive the trajectory data directly from v_s.

For any {time interval, vehicle} pair, multiple trajectory measurements could exist. However, without tampering, each vehicle ultimately can only confirm one of them. Thus, the number of votes for a particular trajectory is simply the number of confirming vehicles in the tree. When enough votes for a particular trajectory are obtained (either a majority or some supermajority, depending on the blockchain model), the trajectory can be added to the blockchain as a permanent record.

4.3 Forwarding

Vehicles will packetize the tree lists of trajectory data and forward the data to nearby vehicles. Prior to transmitting, the forwarding vehicle will update the *hop* and *age* metrics so that each transmission is already updated at receipt. To use network transmission capacity more efficiently, tree lists with similar relevance might be bundled together and transmitted as batches.

Recall that each vehicle in a blockchain knows all members of that blockchain and that vehicles may belong to multiple blockchains. To limit superfluous forwarding, vehicles only forward data involving vehicles from their own blockchain group(s). Any vehicle receiving a packet about a vehicle not in its blockchain(s) would simply not forward it.

4.4 Verification

Trajectories that cannot be directly verified by a vehicle's line-of-sight-sensors must, inevitably, accept data from those with direct knowledge. So, our scheme is designed to facilitate detection of bad data received indirectly. An immediate alteration of a new trajectory generated by a vehicle v_s will not succeed because other vehicles directly verify the trajectory of v_s. The alteration would be easily detected. However, later elements of a tree list, which cannot be directly verified, are a more serious concern.

The existence and sharing of multiple copies of trajectory data, received from multiple sources, provides security against isolated incidents of tampering. For example, if a tree list entry shows a vehicle confirming two different trajectories, vehicles can recognize and exclude the vehicle providing the bad data. Two foreseeable malicious actions are altering trajectory data and altering vehicles in the tree lists:

Altering Trajectory Data: If a vehicle v forwards an incorrect trajectory, other vehicles earlier in the sequence will recognize it. They can see that v sent bad data that did not match what they confirmed. Downstream vehicles from v will likewise recognize that v's claims conflict with others from the same source, the parent node of v in the tree list.

Altering Vehicle Identities: Suppose a vehicle v changes the tree list, either altering a vehicle identity v' or changing the sequence of confirmations. Then, v transmits the changed list. Again, recipients will recognize the change. If they received data directly from v', they know that v changed the data and will not confirm it. This will prevent

v from gaining a majority in the blockchain voting process even if all other vehicles confirm both versions of the tree list.

Malicious alteration attempts like the above by single vehicles are typically easily detected. However, coordinated malicious behavior by groups of nodes cannot always be stopped. Still, this is not a disqualifying weakness for blockchain. In any shared database system, some users must be able to make updates and might act maliciously. Blockchain relies on making coordinated malicious behavior prohibitively difficult – there are simpler and more certain ways to cause harm than getting multiple vehicles to travel together and deliberately manipulate vehicular safety data.

5 Conclusion

This paper presented a framework for combining blockchain and SBSD as a secure, scalable, and decentralized solution to sharing information in IVSAs. Certainly, there remain many details to clarify and questions to resolve. Our next steps will be to clearly model blockchain implementations, including security policies and membership management. Then, we will run the corresponding simulations of V2V applications. This will provide better understanding of the tradeoffs from different blockchain models regarding the timeliness, completeness, and security of IVSA data. Many performance comparisons can be made for variables such as:

- Data delivery and confirmation speed: Assuming trustworthy vehicles, how quickly can data be shared and confirmed within the entire group?
- Faulty or malicious group members: How is IVSA performance affected by isolated or small groups of malicious vehicles?
- Group membership convergence: How is IVSA performance affected by frequent changes in group membership?
- Network overhead: How much network traffic is required to deliver group membership information?
- Reliability: What guarantees are given that required information arrives?

In the coming years, vehicular safety technology will continue to be developed and standardized. Growing adoption of 5G technology will alleviate data transmission bottlenecks, allowing more data to be shared faster and over larger areas. Self-driving cars will become more common and perhaps ubiquitous. We look forward to seeing how these trends converge in vehicular safety applications.

References

1. Consumer Reports Safety First: Car Crashes, Innovation, and Why Federal Policy Should Prioritize Adoption of Existing Technologies to Save Lives, June 29 2020
2. Hussain, R., Lee, J., Zeadally, S.: Trust in VANET: a survey of current solutions and future research opportunities. IEEE Trans. Intell. Transp. Syst. **22**(5), 2553–2571 (2021)
3. Ouksel, A., Lundquist, D.: Demand-driven publish/subscribe in mobile environments. Wireless Netw. **16**(8), 2237–2261 (2010)

4. Federal Communications Commission: FCC 20-164: Use of the 5.850 to 5.925 GHz Band, November 2020
5. Sherman, A., Javani, F., Zhang, H., Golaszewski, E.: On the origins and variations of blockchain technologies. IEEE Secur. Priv. **17**(1), 72–77 (2019)
6. Li, X., Jiang, P., et al.: A survey on the security of blockchain systems. Futur. Gener. Comput. Syst. **107**, 841–853 (2020)
7. Zhang, X., Chen, X.: Data security sharing and storage based on a consortium blockchain in a vehicular ad-hoc network. IEEE Access **7**, 58241–58254 (2019)
8. Arora, S., Kumar, G., Kim, T.: Blockchain based trust model using Tendermint in vehicular adhoc networks. Appl. Sci. **11**(5), 1998 (2021)
9. Gao, F., Zhu, L., et al.: A blockchain-based privacy-preserving payment mechanism for vehicle-to-grid networks. IEEE Netw. **32**(6), 184–192 (2018)
10. Feng, Q., De, H., et al.: BPAS: blockchain-assisted privacy-preserving authentication system for vehicular ad hoc networks. IEEE Trans. Industr. Inf. **16**(6), 4146–4155 (2020)
11. Lin, C., et al.: BCPPA: a blockchain-based conditional privacy-preserving authentication protocol for vehicular ad hoc networks. IEEE Trans. Intell. Transp. Syst. (early access) **22**, 7408–7420 (2020)
12. Ren, Y., Li., X., et al.: Privacy-preserving batch verification signature scheme based on blockchain for vehicular ad-hoc networks. J. Inf. Secur. Appl. **58**, 102698 (2021)
13. Shrestha, R., Nam, S.: Regional Blockchain for vehicular networks to prevent 51% attacks. IEEE Access **7**, 95033–95045 (2019)
14. Joshi, G., Perumal, E., et al.: Toward blockchain-enabled privacy-preserving data transmission in cluster-based vehicular networks. Electronics **9**(9), 1358 (2020)
15. Kim, S.: Impacts of mobility on performance of blockchain in VANET. IEEE Access **7**, 68646–68655 (2019)
16. Xu, B., Ouksel, A., Wolfson, O.: Opportunistic resource exchange in inter-vehicle ad-hoc networks. In: 5th Conference on Mobile Data Management, pp. 4–12, Berkeley, California (2004)