A Transparent Electoral System Technology (ATEST): Increasing trust in elections through an open, immutable ledger.

Sam Williams
sam@arweave.org

Sebastian Campos Groth
sebastian@arweave.org

India Raybould
india@arweave.org

Status: DRAFT
July 21, 2021

Abstract and Executive Summary

Trust in the electoral system is at the core of all functional democratic societies. When the public's trust in the legitimacy of these systems wanes, the social consensus around the entire governance infrastructure becomes threatened. In 2021, across the political spectrum, trust in election integrity is at an all time low.

Decentralized ledgers are a breakthrough technology that can dramatically improve the relationships of trust in a society. This family of technologies allows groups of people to delegate trust to rigorously-proven mathematics and computer networks, rather than to human-operated, and potentially error-prone, institutions. The first deployment of such a technology was in the Bitcoin blockchain [13], an experiment in building a currency that does not depend on trust in a centralized minting authority. Despite the progress of this initial experiment, with Bitcoin's marketcap recently exceeding 1 trillion USD [6], the technology has traditionally been plagued with scalability, [4, 22] security [10, 15], and ecological issues [2].

Iterating on the traditional blockchain structure, Arweave is a scalable and secure permanent information storage ledger [20]. In the Arweave network a decentralized collection of machines is tasked with the long term storage and replication of data. This system uses a sustainable economic endowment structure to ensure resilience and permanence of the ledger's contents. Anyone, anywhere can access information on this ledger without cost, and write to it for a small fee.

In this paper we present ATEST (A Transparent Electoral System Technology), a new technological infrastructure for administering elections that removes the need for trust in human-operated procedures. Instead, this trust is placed in the mathematics that have long powered maintenance of government secrets, e-commerce on the internet, and private communications via email.

1 Introduction

This paper outlines a recommendation for improving election trust through the use of a decentralized, publicly-accessible and -verifiable ledger. The objectives of this system are as follows:

1. Allow voters to validate that their ballot, as they cast it, has been correctly counted during an election. This objective must be achieved without exposing the voter's electoral choice to any third party.

2. Allow citizens and all interested parties to validate the integrity of election voter registration rolls.

3. Make sufficient records of the integrity of all elections permanently and publicly available for permissionless validation by any party in the future.

2 Architecture

At the core of ATEST (A Transparent Electoral System Technology) is a protocol composed of three parts; a private, permissioned Arweave ledger, the public global Arweave network, and a uniform data format representing interactions during the process of an election.

In the proposed architecture a hybrid approach between a public blockchain network and a private, permissioned system is presented. This system captures the benefits of a public network (transparency, auditability, and data permanence) while also providing the control and security of a private environment.

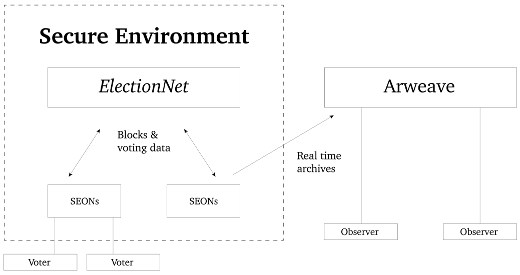

Fig.1: Overview of the ATEST ElectionNet architecture.

2.1 ElectionNet

Upon initialization of election preparation, a new, private, "ElectionNet" network should be created. Ahead of initialization of a new ElectionNet network, a set of valid block-producing nodes must be agreed. While this may differ from deployment to deployment, in the United States, it seems appropriate that State Election Offices should be granted this responsibility. During the initialization process, public keys from these State Election Office Nodes (SEON) must be circulated between participants and included in the first block.

During operation, SEONs will stochastically produce blocks at an appropriate frequency as determined by the deployer. Each block will be signed by the SEON's private key such that all participants can validate its authenticity. Block production can subsequently be performed in a decentralized manner, tolerant to the failure of any one or more nodes in the network. This decentralized block production capacity is similar to traditional Proof-of-Work blockchains (such as Bitcoin [13]) but with two fundamental differences:

1. Block production does not incur significant computational work. This avoids any unnecessary ecological consequences as a result of an electoral undertaking.

2. Consensus in the network is impervious to traditional blockchain security attack vectors. This includes resilience to classic "51%" attacks [15] in Proof-of-Work blockchains, and economic attacks [18] in Proof-of-Stake blockchains.

During operation of the ElectionNet (from voter registration through to finalization) SEONs will receive signed interactions from voters. Upon receipt, SEONs validate these interactions, and include them in bundles (using Arweave's network standard ANS-104 [3] technology), which are subsequently incorporated into blocks. In a traditional blockchain network, all validating nodes must have local access to every interaction included in a block. This creates a fundamental bottleneck to interaction processing speed – that of the bandwidth of the poorest-performing node. Arweave's bundling technology circumvents this limitation by allowing bundle producers to publish only a Merkle root (MR) [11, 21] of an essentially arbitrarily large collection of interactions in a block (limited at $2^{256} - 1$ bytes per bundle). This approach is significantly more efficient than if they were required to publish the full interaction data. Validating nodes then come to consensus about just this MR during block production, rather than about every interaction the block contains, allowing for delayed synchronization of the underlying data, and avoiding the aforementioned bottleneck in interaction processing speed. This will allow ATEST to trivially surpass the required interaction processing speed for an election amongst even billions of voters.

2.2 Public ledger synchronization

Periodically throughout the lifetime of an ElectionNet, SEONs will submit bundled copies of new blocks and their contained interactions to the public Arweave ledger for inclusion in blocks produced by the ledger's Proof-of-Access [19] storage system. In doing so, each SEON will give an attestation of its perspective of the state of the ElectionNet with immutable timestamps – allowing permissionless verification of the metadata globally, in real time. For example, this system would allow election observers in third party nations to validate that the state of the ElectionNet from the perspective of the Maine SEON is the same as that of the Wisconsin SEON – or any other pairing.

These additional security factors provided by the data being uploaded to the Arweave network will greatly increase verifiability and trust in election results while also ensuring that valuable records are kept for generations to come, backed by Arweave's sustainable economic endowment mechanism [20].

2.3 Voter identity management

In order to allow cryptographically secure voting, a 1:1 mapping between public-private key pairs and registered voters is required. In general, the proposed architecture is flexible to different key types and generation methods – the form of asymmetric public-private cryptography may be altered as appropriate for the given deployment. Choices compliant with the Commercial National Security Algorithm Suite [16] (CNSA Suite) would include the Elliptic Curve Digital Signature Algorithm [9] (ECDSA), as well as RSA [14] with a minimum key length of 3072 bits. Similarly, vehicles for holding users' private key data may vary according to implementation specifics and requirements, however they should typically reside in some form of hardware security module (HSM).

2.4 Voter registration process

After the establishment of the ElectionNet in preparation for a voting process, each SEON will allow users to begin registration of their public keys and identities. Each SEON controlled by its own state's electoral system may have different requirements surrounding voter registration. Subsequently, verification of voter registration will not be performed by other nodes in the ElectionNet, instead simple attestation from the SEON (signed by its private key) will be produced and transmitted to the other nodes in the network for inclusion in the block. This provides flexibility, allowing each state to implement voter registration checks as required by law, while also enforcing high levels of accountability and providing auditability to the process of adding voters to the rolls. For example, through the use of this system observers will be able to permis-

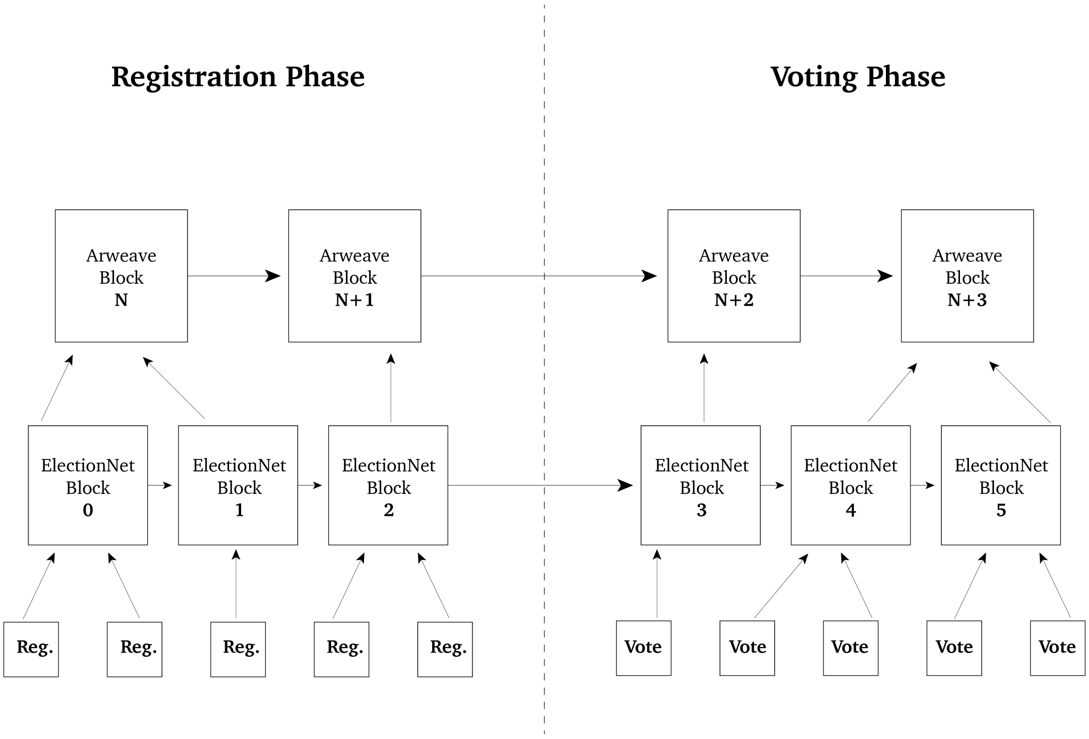

Fig.2: The proposed ATEST block production protocol.

sionlessly validate metadata surrounding submission of new voters to the rolls – including the frequency of data submissions, the locations of new voters added to the system, and other appropriate interactions. If the system had been in place during prior election cycles it would have been possible to simply dispel concerns in a number of states that large numbers of voters were falsely added to the rolls shortly before the elections took place. With the proposed system such discrepancies would now be visible and detectable in real time through statistical analysis of the speed and timing of the voter registrations.

2.5 Ballot casting and counting procedure

After the voter registration period has closed the ballot casting and counting process may begin. At this point, each SEON will begin to receive signed votes from users which they will bundle and submit to the ElectionNet. As with all other submissions to the ElectionNet, this data will be inserted into the private ledger (written into blocks), and those blocks will be archived into the public Arweave ledger. Just as with voter registration information, this process will allow observers to watch the pace of votes, and other metadata, in order to validate that votes cast in the election stem from legitimately registered voters. This will create an extremely strong mechanism for dispelling traditional election integrity concerns surrounding "ballot stuffing".

This specification does not mandate any specific vote-counting privacy technologies are used, instead it provides broad support for a range of different technological solutions. One such compatible privacy scheme would be a simple commit-reveal mechanism in which voters submit a hash of their preference along with a nonce to their SEON. The SEON would then keep the nonce and preference data private, but publish the output hash, its associated metadata, and the vote's signature into its data bundles in the ElectionNet. After the election is over and the votes are to be counted, the SEON may privately reveal the voting data to the appropriate parties who can tally and officiate the results of the election, validating the given data against the tamper-proof timestamps and cryptographically-signed votes provided in the ElectionNet and on the public Arweave ledger.

Another such privacy-preserving vote-tallying mechanism could be created through the use of fully homomorphic encryption [1, 7, 17, 5] (FHE) schemes. While enforcing more rigorous requirements upon encryption algorithm choices, an FHE scheme would also allow users additional validation options. For example, with such a scheme it is possible for users to validate with cryptographic confidence that their particular vote was correctly included in the tallying of the final counts, without any other party being able to view information about their voting choice. Both Google [8] and Microsoft [12] have recently released FHE libraries that would be compatible with the described ATEST specification in this paper. It is notable, however, that FHE schemes have only recently reached maturity and therefore potentially present a more significant danger of cryptographic breaks and other issues. By contrast, the simpler 'commit-reveal'-style schemes rely only upon more established cryptographic primitives, at the cost of additional features.

3 Future Work

The core Arweave team will be spearheading the development of this proposed architecture, after collating and integrating feedback from potential stakeholders.

The backend architecture and framework will be published by the core Arweave team as open source software, ensuring that it can be used by anyone for any purpose without impediment. As a consequence, users of the system will incur no software development or licensing costs relating to the core architecture. Deployers of the system will likely, however, want to build appropriate frontend user interfaces and ballot tabulation mechanisms customised for their deployment, which would incur the traditional financial costs associated with such design and development work.

4 Conclusion

This paper outlines a fully open source architecture for increasing verifiability and public trust in election systems, based on decentralized ledger technologies. The solution has high security guarantees (circumnavigating all typical decentralized ledger attack vectors), plus high scalability and verifiability, with low deployment costs and lead time. Feedback and suggestions from any potential stakeholders or users are strongly encouraged and welcomed by the core Arweave team.

Points of contact

Sam Williams – sam@arweave.org

Sebastian Campos Groth – sebastian@arweave.org

India Raybould – india@arweave.org

References

[1] Frederik Armknecht et al. *A Guide to Fully Homomorphic Encryption.* Cryptology ePrint Archive, Report 2015/1192. `https://eprint.iacr.org/2015/1192`. 2015.

[2] Liana Badea and Mariana Claudia Mungiu-Pupzan. “The Economic and Environmental Impact of Bitcoin”. In: *IEEE Access* 9 (2021), pp. 48091–48104.

[3] Joshua Benaron. *Arweave Network Standard 104.* URL: `https://github.com/joshbenaron/arweave-standards/blob/ans104/ans/ANS-104.md`.

[4] Anamika Chauhan et al. “Blockchain and scalability”. In: *2018 IEEE International Conference on Software Quality, Reliability and Security Companion (QRS-C).* IEEE. 2018, pp. 122–128.

[5] Jung Hee Cheon et al. "Homomorphic encryption for arithmetic of approximate numbers". In: *International Conference on the Theory and Application of Cryptology and Information Security*. Springer. 2017, pp. 409–437.

[6] CoinGecko. *Bitcoin price chart*. URL: `https://www.coingecko.com/en/coins/bitcoin`. (accessed: 21.07.2021).

[7] Craig Gentry. "Fully homomorphic encryption using ideal lattices". In: *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*. 2009, pp. 169–178.

[8] Google. *Google Fully Homomorphic Encryption*. URL: `https://github.com/google/fully-homomorphic-encryption`.

[9] Don Johnson, Alfred Menezes, and Scott Vanstone. "The elliptic curve digital signature algorithm (ECDSA)". In: *International Journal of Information Security* 1.1 (2001), pp. 36–63.

[10] Ghassan Karame. "On the security and scalability of Bitcoin's blockchain". In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. 2016, pp. 1861–1862.

[11] Ralph C Merkle. "A digital signature based on a conventional encryption function". In: *Conference on the Theory and Application of Cryptographic Techniques*. Springer. 1987, pp. 369–378.

[12] Microsoft. *Microsoft SEAL*. URL: `https://github.com/microsoft/SEAL`.

[13] Satoshi Nakamoto. "Bitcoin: A peer-to-peer electronic cash system". In: *Decentralized Business Review* (2008), p. 21260.

[14] Ronald L Rivest, Adi Shamir, and Leonard Adleman. "A method for obtaining digital signatures and public-key cryptosystems". In: *Communications of the ACM* 21.2 (1978), pp. 120–126.

[15] Sarwar Sayeed and Hector Marco-Gisbert. "Assessing blockchain consensus and security mechanisms against the 51% attack". In: *Applied Sciences* 9.9 (2019), p. 1788.

[16] National Security Agency Central Security Service. *Commercial National Security Algorithm Suite*. URL: `https://apps.nsa.gov/iaarchive/programs/iad-initiatives/cnsa-suite.cfm`. (accessed: 16.07.2021).

[17] Marten Van Dijk et al. "Fully homomorphic encryption over the integers". In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2010, pp. 24–43.

[18] B. Victor. "Attack vectors in proof of stake blockchains". 2018. URL: `http://arks.princeton.edu/ark:/88435/dsp014j03d238m`.

[19] Sam Williams and Lev Berman. *Arweave Network Standard 103*. URL: `https://github.com/ArweaveTeam/arweave-standards/blob/ans-103/ans/ANS-103.md`.

[20] Sam Williams et al. *Arweave: A Protocol for Economically Sustainable Information Permanence*. Tech. rep. 2019. URL: `https://www.arweave.org/yellow-paper.pdf`.

[21] Xiaojing Yang, Jinshan Liu, and Xiaohe Li. "Research and Analysis of Blockchain Data". In: *Journal of Physics: Conference Series*. Vol. 1237. 2. IOP Publishing. 2019, p. 022084.

[22] Qiheng Zhou et al. "Solutions to scalability of blockchain: A survey". In: *IEEE Access* 8 (2020), pp. 16440–16455.