

CHAPTER 22

Blockchain Electronic Vote

Pierre Noizat

Co-founder of Paymium, Bitcoin Exchange and provider of e-commerce solutions, Paris, France

Contents

22.1	The Problem with Proprietary Voting Systems	453
22.2	Open-Source, Free Software Electronic Transaction and Voting Systems	453
22.2.1	Preelection: Know the list of candidates	454
22.2.2	Check the list of voters	454
22.2.3	Prepare the ballot	456
22.2.4	Counting votes on election day	459
22.2.5	Postelection: Showing independently verifiable results	459
22.3	Conclusion	460
	References	461

22.1 THE PROBLEM WITH PROPRIETARY VOTING SYSTEMS

Existing electronic voting systems all suffer from a serious design flaw: They are proprietary, that is, centralized by design, meaning there is a single supplier that controls the code base, the database, and the system outputs and supplies the monitoring tools at the same time.

The lack of an open-source, independently verifiable output makes it difficult for such centralized systems to acquire the trustworthiness required by voters and election organizers.

This design flaw is therefore limiting electronic vote applications at a time when growing computer literacy and usage should in fact foster their widespread adoption.

Without an easy, free access to effective, secure electronic voting technology, political participation is limited to on-site elections that are few and far between, given the high setup cost of election preparation, supervision, and postelection operations.

Electronic voting is not meant to replace traditional elections but can provide a much-needed complementary voting method.

22.2 OPEN-SOURCE, FREE SOFTWARE ELECTRONIC TRANSACTION AND VOTING SYSTEMS

Since the invention of the Bitcoin protocol (Nakamoto, 2008), the Bitcoin network has grown to become today the most powerful distributed computing network on Earth,

providing a very secure, free software, database infrastructure to store electronic transaction data of any kind.

The complete Bitcoin transaction database is referred to as the blockchain because bitcoin transaction records are paginated in blocks appended to the database every 10 min on average.

An electronic vote is essentially an electronic transaction whereby a voter, given some voting credits, will spend them in favor of one or more candidate recipients.

Candidate recipients can be people like in a presidential election or options to choose from like in a referendum election.

This chapter describes how to leverage the availability of the Bitcoin blockchain as a secure transaction database, to log votes and audit vote results. The proposed method involves Merkle trees (Merkle, 1980) for voters' list verification and block explorers for vote count check.

The solution protects the privacy of the vote by using a simple form of Bitcoin pay-to-script-hash scripts (Andresen, 2012) (<https://github.com/bitcoin/bips/blob/master/bip-0016.mediawiki>) with 2-of-3 multisignature addresses.

I have developed a demo voting application based on my proposal to assess usability and performance, but testing has only begun and is far from over: I am sharing here early results to get the discussion going. The application is deployed on the web site: www.unchain.voting.

Here is how the whole process unfolds in five steps.

22.2.1 Preelection: Know the list of candidates

Each candidate (C) assigns a public key $KeyC$ to the voter and publishes everywhere a Bitcoin address like 1Martin., that is, a vanity address (<https://en.bitcoin.it/wiki/Vanitygen>) prepared by the organizers for candidate Martin.

For the secrecy of each ballot, each $KeyC$ is different for each voter, so the candidate must generate n public keys $KeyC$ for n voters.

Each candidate should be required by the organizers to publish the Merkle root of a hash tree comprising all of his public keys $KeyC$.

Alternatively, the voting application can generate the keys $KeyC$ on behalf of the candidate, although this option could be avoided to reduce the need to trust the organizers.

Using the voting application, each voter should be able to check her branch of the hash tree, linking the key $KeyC$ to the root.

22.2.2 Check the list of voters

Each voter (B) is also assigned by the election organizers (e.g., an association A) a public key $KeyA$ and a nonce value that can be used by the voter as a secret access key to log on to her account.

These access keys will be sent securely by the organizers to each voter.

The voters' list can be represented in a Merkle tree (Figure 22.1): The application gives each voter the list of interior hashes, linking her public key KeyA to the Merkle root.

With the Merkle tree published by the application and with the standard computation formulas below, the voters' list can be independently verified by all parties because every voter can check the path from her leaf node to the Merkle root and make sure that the number of leaf nodes is equal to the number of voters.

This is how interior hashes are calculated as 64-character string, using Ruby scripting:
`sha256_base64 "#{sum_of_upper_left_and_right_nodes}|#{upper_left_hash}|#{upper_right_hash}"`

Likewise, leaf node hashes are calculated as follows:

`sha256_base64 "#{voter_name}|#{voter_credit_balance}|#{voter_nonce}"`

The voter credit balance can be initialized to one in the simple case of a presidential election or to n if the voter can choose n of m candidates like in the election of n board members. Voter credit is updated when the voter casts her vote, spending all or part of her voting credit.

`sha256_base64` is the standard method for generating the 64-character hash of a string:
`sha256_base64(string) = OpenSSL::Digest::SHA256.new.digest(string).unpack('H*').first`

The path to the root for one of the voters in the above example tree would look like this:

Voter name: stephan

Voter credit balance: 0.01 ₿

Account Nonce (voter access key): 958a78b68c8f53f89c3433c1fde98152

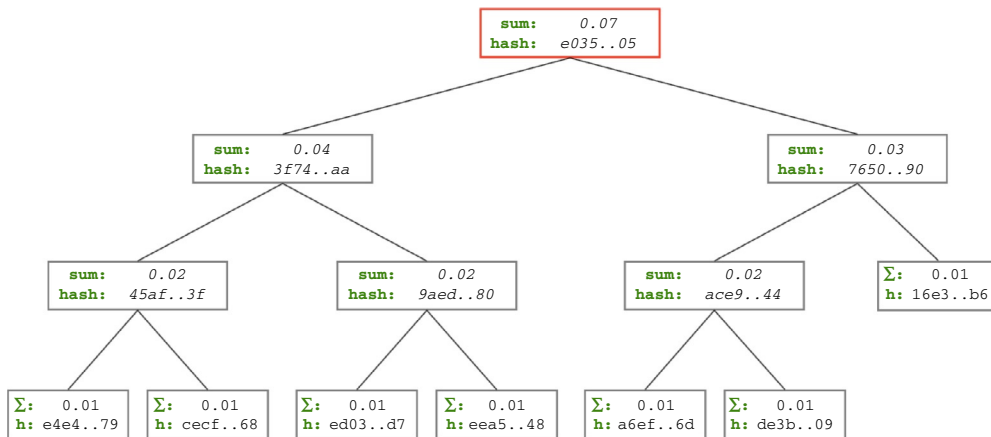


Figure 22.1 Merkle tree of a seven-voter list.

Account Hash (leaf node hash): eea58cf19ed83072a9ea650e1c71a100f8f22129196cf52d1d91c4ecf68bdc48

Interior Hash #1:

9aed30f6447d11e852873d4896d43a13152773b7ba2dae65b1eb342854c9de80

Interior Hash #2:

3f74cef73f11bbc2eca69e70b985f7a3e6a69d292f4dc0fb9d0fbcdd5453dfaa

Interior Hash #3 (Merkle Root):

e035cfb7615443b4e49bc1a0bd4b16bdb28ec79ad8c409f5bbf55dea9cbf7f05

As a precautionary option, the public key *KeyA* should be random and should not be on the Bitcoin elliptic curve bitcoin, for instance, using a side chain, to prevent the organizers (A) from producing a valid signature.

This precaution is also independently verifiable and reduces the need to trust the organizers. However, the complexity of using a side chain instead of the blockchain should not be underestimated.

Additionally, the application assigns a unique public key *KeyB* to the voter.

As an option, to further reduce the need to trust the organizers, *KeyB* can be generated independently by the voter herself. However, opening this option would entail some voter education, support efforts, and less usability.

As usual, the choice for an optimal convenience/security trade-off should be guided by common sense and knowledge of the voters' demographics, including computer literacy. Generally speaking, organizers should not expect that people will be able to handle private keys themselves in a secure manner.

22.2.3 Prepare the ballot

Before election day, with the A, B, and C public keys recorded by the application, the application can create a *2-of-3 multisignature address* and a transaction sending to it a bitcoin micropayment (about the price of a postage stamp) upon the voter's confirmation of her vote.

The voter's ballot now has a cryptographically secure representation in this transaction, which is broadcast to the bitcoin network nodes by the application.

At this point, it should be noted that the voter can check independently and immediately that her vote has been submitted to the bitcoin network by looking up the multisignature address with any block explorer such as blockchain.info or blockr.io: Although unconfirmed yet, the transaction is showing 2–3 s at most after she confirmed her vote.

After a few hours, the ballot is securely logged on the blockchain with an increasing number of network confirmations.

The output script (destination) of her micropayment transaction looks like this:

OP_HASH160 Hash OP_EQUAL

Where *Hash* is the hash calculated by her bitcoin wallet from the following (*SerializedScript*) script:

```
OP_0 OP_2 KeyA KeyB KeyC OP_3 OP_CHECKMULTISIG
```

This transaction is compliant with BIP-16.

Each key is a concatenation of the x and y coordinates of a point. If the point is on the bitcoin Koblitz curve, then it is a Bitcoin public key and there is a corresponding private key. Coins sent to the multisignature address can be spent only if at least two of the keys are Bitcoin public keys. A key is represented as a 128-hexadecimal character string, each coordinate being represented as 32 bytes.

Here is an example:

KeyA:

```
7d56d13563243998e4934e189574db4eca3b1ac1e0f20f8e136ec047a202151da60a38
a561bc4debd28c0433c51b2f21f1df58c53d4e69c828351dff4ce1eca
```

KeyB:

```
d3363437960978fef861678d9ae0e18544cf999892c0375e7b2856a5258d740f0dfd31
d3b1547b0847c1e65f4d7228a2e7223cc18e42b4271bf7421a54f03d48
```

KeyC:

```
2fdede700db4d7c5e40cbe90d61961240e6fab869270c21d024614b83803adfc84e001
d313a24910317ec3a7e73c6dcdcf828acb639d8048cc4efed95b77c8eb9
```

The following is the multisignature address whose funding represents the vote of B in favor of C: 3NSxmhDPpYak4HH5BYJdsNDq5s1bE121oH.

Because this address was funded, it can be found in a transaction output of the blockchain by anyone using a block explorer. More precisely, the address is derived from the 20-byte hash of the redeem script of the transaction recorded in the blockchain. The redeem script itself, a simple concatenation of the three public keys, will not be visible until the output is spent on a subsequent transaction, after election day.

The application displays a list of all the funded multisignature addresses as an indicator of the vote participation rate.

There is no way to guess neither the voter (B) nor the candidate (C) from a multisignature address without knowing all three public keys (KeyA, KeyB, and KeyC) and knowing to whom they belong.

In particular, the candidates have no way of knowing their rankings at this point.

A voter can check if the multisignature address representing her vote appears in the list.

In the demo application, organizers generate and fund a separate address for each voter in the preelection period. This address (Address B) is derived from KeyB to keep things simple.

On election day, the unspent output of Address B can be used by the application to fund the multisignature address corresponding to the voter's choice.

If multiple choices are offered in the election like in a board member election, several multisignature addresses can be linked to the voter's account to reflect her selection, each address representing a selected candidate or chosen option.

Since we want the multisignature address funding transactions to show up in the blockchain, preferably before the vote count, organizers must consider the settings of transaction fees. If transaction fees were set too low, a funding transaction could end up being left in limbo, stuck in the memory pool of the miners, and ultimately dropped.

According to the Bitcoin Wiki (https://en.bitcoin.it/wiki/Transaction_fees), a transaction may be sent without fees if the following conditions are met:

- It is smaller than 1000 bytes.
- All outputs are 0.01 BTC or larger.
- Its priority is large enough, where

$\text{priority} = \text{sum}(\text{input_value_in_base_units} * \text{input_age}) / \text{size_in_bytes}$

The typical size of a transaction from one standard address to one multisignature address is around 223 bytes.

As of core client version 0.3.21, *transactions needed to have a priority above 57,600,000 to avoid the enforced limit. This threshold is written in the core client code as $\text{COIN} * 144 / 250$, suggesting that the threshold represents a one day old, 1 btc coin (144 is the expected number of blocks per day) and a transaction size of 250 bytes.*

*So, for example, a transaction that has 2 inputs, one of 5 btc with 10 confirmations, and one of 2 btc with 3 confirmations, and has a size of 500bytes, will have a priority of $(500000000 * 10 + 200000000 * 3) / 500 = 11,200,000$*

Therefore, the recommended setup process to avoid paying fees to the miners is to use 0.01 BTC as the transaction amount for each voter and to fund Address B at least 1 day before election day.

The challenge is about paying minimal mining fees while ensuring that the transactions will be confirmed.

It should be anticipated that, in the future, miners' fees will evolve from rule-driven fees to market-driven fees, turning the above fee-avoiding recommendation into a fee-minimizing strategy.

Because the time period when the coins are needed equals roughly 2 days, the amount of the transaction should be scaled according to the number of voters: As the amount is driven lower by a large number of voters, mining fees will start to kick in. As a result, organizers may be able to find an optimal trade-off combining financial costs and mining fees.

In fact, as noted in the description of [Section 22.2.3](#), there is no strong requirement for a fast confirmation time as long as the voter can see that her vote transaction has been broadcast to the network nodes.

Confirmations are needed only to etch the transaction in the blockchain forever and to make verifications possible at any time in the future.

If large-scale operations like a national election are on the agenda, the organizers should consider mining themselves to ensure that the vote transactions will eventually make their way to the blockchain: the more hashing power harnessed by the organizers, the lower the statistical cap on the maximum confirmation time.

22.2.4 Counting votes on election day

At the end of election day, the application can simply display the compiled results from its database.

Compared to a traditional, proprietary system, the added feature is the ability to display a funded multisignature address for each vote record.

The blockchain does not say much about this address yet until its coins are spent to another address.

The election results will be made independently verifiable in the next and final step.

22.2.5 Postelection: Showing independently verifiable results

Now is the time to show the world a lasting, easily verifiable proof that the vote count was not rigged in any way.

Let's say one of the candidate is Martin: Martin has advertised a special bitcoin vanity address 1Martin.. on social networks during the election campaign.

For each voter, the application uses the two private keys corresponding to KeyB and KeyC, respectively, to sign a transaction, spending the coins from the 2-of-3 multisignature address to the address 1Martin., unequivocally linking the vote to the candidate.

The transaction from the multisignature address to 1Martin.. includes the following input:

SignatureB SignatureC SerializedScript,

where SerializedScript designates the redeem script converted from its binary form to a hexadecimal string.

Looking up the multisignature address in the blockchain using any block explorer, it is easy to find the redeem script in the spending transaction, in other words the serialized input script of the transaction sending coins from 3NSxmh... , revealing the three public keys.

An observer can use the Merkle tree published by the candidate before the election to check the validity of the interior hashes connecting KeyC to the Merkle root. Similar verifications applied to KeyA and KeyB will allow the observer to assert the validity of the multisignature address.

A discrepancy between the vote count provided by the application on election day and the blockchain transactions could be spotted easily with any block explorer software.

However, it is impossible for an observer to link a multisignature address to a voter insofar as KeyB cannot be linked to the voter identity by the observer.

Incidentally, this system uses an advantage of a digital signature compared with a manual signature: a digital signature identifies its author only with her consent, if and when she reveals her identity.

The votes, while preserving the secrecy of the ballots, are perfectly verifiable by candidates and voters, independently of any organization.

To complete the postelection operations, the organizers collect the funds from address 1Martin.. and from the other candidates' vanity addresses that were prepared in the pre-election period.

Would it be simpler to just create a new blockchain and wallet for each election cycle?

It is in fact simpler to use an existing blockchain.

Besides, no other blockchain comes close to the Bitcoin blockchain in terms of protection against a 51% attack led by one of the candidates.

Creating a new blockchain dedicated to a particular vote amounts to requesting the voters to trust more the organizers and all the candidates: the organizers or some candidate could take control of the new blockchain, mine some votes, and drop some other votes. This would seriously disrupt the verification of the vote count.

22.3 CONCLUSION

Before the invention of Bitcoin, electronic voting solutions were not satisfactory because they were not easily auditable and not sufficiently transparent neither for candidates nor for voters. In addition, they require a costly, labor-intensive setup.

With the Bitcoin blockchain, any community can organize a free, secure electronic voting.

Initially, communities, associations, or listed corporations should consider using this technology for local elections, board elections, or general assembly voting of their members or shareholders.

Scaling up to national elections might require direct involvement of the organizers in mining operations or some level of cooperation with miners.

The proposed system does not solve all the issues associated with electronic voting, but it does provide a valuable alternative to current, proprietary electronic voting systems with the following benefits:

- Free, open-source peer-reviewed software
- Ubiquitous
- Secure
- Protecting the secrecy of the ballots
- Allowing free, independent audits of the results
- Minimizing the trust level required from the organizers

REFERENCES

- Andresen, G., 2012. BIP-0016 (Bitcoin Improvement Proposal), January 2012.
- Merkle, R.C., 1980. Protocols for public key cryptosystems. In: *Proceedings of the 1980 Symposium on Security and Privacy*, April 1980. IEEE Computer Society, pp. 122–133.
- Nakamoto, S., 2008. Bitcoin: a peer-to-peer electronic cash system, November 2008.