

Civitas: Toward a Secure Voting System

Michael R. Clarkson Stephen Chong Andrew C. Myers

Department of Computer Science, Cornell University

{clarkson, schong, andru}@cs.cornell.edu

Abstract

Civitas is the first electronic voting system that is coercion-resistant, universally and voter verifiable, and suitable for remote voting. This paper describes the design and implementation of Civitas. Assurance is established in the design through security proofs, and in the implementation through information-flow security analysis. Experimental results give a quantitative evaluation of the tradeoffs between time, cost, and security.

1. Introduction

Electronic voting is now a reality—and so are the many errors and vulnerabilities in commercial electronic voting systems [2, 8, 48, 73]. Voting systems are hard to make trustworthy because they have strong, conflicting security requirements:

- *Integrity* of election results must be assured so that all voters are convinced that votes are counted correctly. Any attempt to corrupt the integrity of an election must be detected and correctly attributed.
- *Confidentiality* of votes must be assured to protect voters' privacy, to prevent selling of votes, and to defend voters from coercion.

Integrity is easy to obtain through a public show of hands, but this destroys confidentiality. Confidentiality can be obtained by secret ballots, but this fails to assure integrity. Be-

This work was supported by the Department of the Navy, Office of Naval Research, ONR Grant N00014-01-1-0968; Air Force Office of Scientific Research, Air Force Materiel Command, USAF, grant number F9550-06-0019; National Science Foundation grants 0208642, 0133302, 0430161, and CCF-0424422 (TRUST); and a grant from Intel Corporation. Michael Clarkson was supported by a National Science Foundation Graduate Research Fellowship and an Intel PhD Fellowship; Andrew Myers was supported by an Alfred P. Sloan Research Fellowship. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either express or implied, of these organizations or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

cause of the civic importance of elections, violations of these requirements can have dramatic consequences.

Many security experts have been skeptical about electronic voting [25, 28, 41, 53, 63], arguing that assurance in electronic voting systems is too hard to obtain and that their deployment creates unacceptable risks. Our work, however, was inspired by the possibility that electronic voting systems could be more trustworthy than their non-electronic predecessors. This paper describes and evaluates Civitas, the prototype system we built to explore that possibility. Although not yet suitable for deployment in national elections, Civitas enforces *verifiability* (an integrity property) and *coercion resistance* [45] (a confidentiality property). Civitas does not rely on trusted supervision of polling places, making it a *remote* voting system.

To obtain assurance in the security of Civitas, we employed principled techniques:

- **Security proofs.** The design of Civitas refines a cryptographic voting scheme¹ due to Juels, Catalano, and Jakobsson [45], who proved their scheme secure; we extend the proof to account for our changes.
- **Secure information flow.** The implementation of Civitas is in Jif [54, 56], a language which enforces information-flow security policies.

This validation of the design and implementation supports our argument that Civitas is secure.

The security provided by Civitas is not free. Tradeoffs exist between the level of security provided by Civitas tabulation, the time required for tabulation, and the monetary cost of tabulation. To better understand these tradeoffs, we studied the performance of Civitas. The results reveal that (with reasonable security and time parameters), the marginal cost of tabulation is as low as 4¢ per voter. Since the current cost of a government election in a stable Western democracy is \$1 to \$3 per voter [38], Civitas can provide increased security at little additional cost.

¹For clarity, we define *voting systems* as implementations, *voting schemes* as cryptographic protocols, and *voting methods* as algorithms that aggregate voters' preferences to produce a collective decision.

Developing Civitas led to several contributions:

- A provably secure voter registration protocol, which distributes trust over a set of registration authorities.
- A scalable design for vote storage that ensures integrity without expensive fault tolerance mechanisms.
- A performance study demonstrating the scalability of secure tabulation.
- A coercion-resistant construction for implementing a ranked voting method.
- A concrete, publicly available specification of the cryptographic protocols required to implement a coercion-resistant, verifiable, remote voting scheme. This specification leverages many results in the cryptographic and voting literature.

Moreover, Civitas is the first voting system to implement a scheme proved to satisfy coercion resistance and verifiability. Thus, Civitas takes an important step toward bringing secure electronic voting to reality.

We proceed as follows. Section 2 discusses the Civitas security model. The design of Civitas is presented in Section 3. Section 4 evaluates the security of Civitas. The implementation of cryptographic components is described in Section 5, and the scalability of tabulation is analyzed in Section 6. The Jif implementation is described in Section 7. Section 8 presents our performance study. Related work is reviewed in Section 9, and some remaining challenges are identified in Section 10. Section 11 concludes.

2. Security Model

The Civitas security model comprises the environment in which Civitas is used, the security properties we require Civitas to satisfy, and the capabilities we ascribe to the *adversary* attempting to subvert those properties.

Remote voting. Electronic voting systems are often designed for *supervised* voting, which assumes trusted human supervision of the voters, procedures, hardware, and software in polling places. But this contradicts society’s trend toward enabling interactions from anywhere at any time. For example, voters in the state of Oregon now vote only by postal mail, and all states receive a substantial fraction—enough to change the outcome of many elections—of their ballots by mail as absentee ballots. As another example, Internet voting is increasingly used by groups such as Debian, the ACM, and the IEEE. Estonia even conducts legally binding national elections using the Internet.

Postal voting and Internet voting are instances of remote voting, which does not assume trusted supervision of polling places. Remote voting is thus a more general problem, and a harder problem, than supervised voting. Because of the evident interest in remote voting, we believe that remote voting

is the right problem to solve. One of our goals was therefore to strike a reasonable compromise between enabling remote voting and guaranteeing strong security properties. This compromise led to two requirements. First, in some circumstances, voters must register at least partly in person. Second, voters must trust the computational device they use to submit votes—though unlike conventional supervised voting, in which voters must trust the particular device supplied by their local election authorities, Civitas enables each voter to choose a supplier and device. We discuss these requirements in Section 4.

Security properties. To fulfill the integrity requirement of Section 1, we require Civitas to satisfy:

Verifiability. *The final tally is verifiably correct. Each voter can check that their own vote is included in the tally (voter verifiability). Anyone can check that all votes cast are counted, that only authorized votes are counted, and that no votes are changed during counting (universal verifiability).*²

We define “verifiability” informally for simplicity, but Civitas satisfies the formal definition given by Juels et al. [45].³

Verifiability improves upon the integrity properties commonly offered by real-world voting systems. For example, real-world systems rarely allow individual voters to verify that their own votes were included in the tally, or to verify the tally themselves. As another example, the commercial electronic voting systems currently deployed in California offer no guarantees that votes are counted correctly [73].

To fulfill the confidentiality requirement of Section 1, a voting system might guarantee *anonymity*, meaning that the information released by the system never reveals how a voter voted. However, for remote voting, anonymity is too weak. Voters might gain additional information during voting that could enable the buying and selling of votes. Such information could also be used to coerce voters. In remote voting, the coercer could even be the voter’s employer or domestic partner, physically present with the voter and controlling the entire voting process. Against such coercers, it is necessary to ensure that voters can appear to comply with any behavior demanded of them. Further, confidentiality must be maintained even when voters collude with the adversary.

Thus, for confidentiality, we require Civitas to satisfy:

Coercion Resistance. *Voters cannot prove whether or how they voted, even if they can interact with the adversary while voting.*⁴

²Universal verifiability was originally defined by Sako and Kilian [66].

³Verifiability could be formulated as the correctness property of secure multi-party computation [33]. Intuitively, this requires that no adversary can change the results of tabulation to be different than if all votes were announced and tabulated publicly.

⁴Removing interaction with the adversary results in *receipt-freeness*, a weaker property originally defined by Benaloh [6].

We define “coercion resistance” informally⁵ for simplicity, but Civitas again satisfies the formal definition given by Juels et al. [45].⁶ This formal definition requires Civitas to defend against attacks in which the adversary demands secrets known to the voter, and attacks in which the adversary demands that the voter submits a value chosen by the adversary. This value might be a legitimate vote or a random value. The adversary may even demand that the voter abstain by submitting no value at all.⁷

A third security requirement that could be added is *availability* of the voting system and tabulation results. Although this would be essential for a national voting system, we do not require our prototype to satisfy any availability property. Some aspects of availability, such as fault tolerance, could be addressed by well-known techniques. Other aspects, such as defending against selective denial-of-service attacks intended to disenfranchise particular groups of voters, are open problems.

Threat model. We require Civitas to be secure with respect to an adversary (essentially due to Juels et al. [45]) with the following capabilities:

- The adversary may corrupt a threshold (made precise in Section 4) of the *election authorities*, mutually distrusting agents who conduct an election. Agents might be humans, organizations, or software components.
- The adversary may coerce voters, demand their secrets, and demand any behavior of them—remotely or in the physical presence of voters. But the adversary may not control a voter throughout an entire election, otherwise the voter could never register or vote.
- The adversary may control all public channels on the network. However, we also assume the existence of some *anonymous* channels, on which the adversary cannot identify the sender, and some *untappable* channels, which the adversary cannot use at all.⁸
- The adversary may perform any polynomial-time computation.

3. Design

Civitas refines and implements a voting scheme, which we refer to as *JCJ*, developed by Juels, Catalano, and Jakobs-




Figure 1. Civitas architecture

son [45]. The differences between our design and JCJ are discussed in Section 9.

3.1. Agents

There are five kinds of agents in the Civitas voting scheme: a supervisor, a registrar, voters, registration tellers, and tabulation tellers. Some of these are depicted in Figure 1. The agents other than voters are election authorities:

- The *supervisor* administers an election. This includes specifying the ballot design and the tellers, and starting and stopping the election.
- The *registrar* authorizes voters.
- *Registration tellers* generate the credentials that voters use to cast their votes.
- *Tabulation tellers* tally votes.

These agents use an underlying *log* service that implements publicly readable, insert-only storage. Integrity of messages in a log is ensured by digital signatures. Agents may sign messages they insert, ensuring that the log service cannot forge new messages. The log service must sign its responses to reads, ensuring that attempts to present different views of log contents to different readers can be detected. Multiple instances of the log service are used in a single election. One instance, called the *bulletin board*, is used by election authorities to record all the information needed for verifiability of the election. The remaining instances, called *ballot boxes*, are used by voters to cast their votes.⁹

3.2. Setup phase

First, the supervisor creates the election by posting the ballot design on an empty bulletin board. The supervisor also identifies the tellers by posting their individual public keys.¹⁰

⁵“Coercion resistance” is used informally throughout the literature. Juels et al. [45] and Delaune et al. [22] give formal definitions in the computational and symbolic models, respectively, of cryptography. The informal definition given above is consistent with both.

⁶Coercion resistance could be formulated as the privacy property of secure multi-party computation. Intuitively, this requires that no adversary can learn any more about votes than is revealed by the results of tabulation.

⁷Note that the requirement to defend voters from forced-abstinence attacks is incompatible with a public record of who has voted.

⁸An untappable channel must provide perfect secrecy, perhaps by being physically untappable or by implementing a one-time pad.

⁹In our prototype, the log service instances are centralized systems provided by the election authorities—the bulletin board by the supervisor, and one ballot box by each tabulation teller. But instances could be made distributed systems to improve availability, and instances could be provided by agents other than the election authorities.

¹⁰A real-world deployment of Civitas would need a public-key infrastructure to certify keys.

Second, the registrar posts the *electoral roll*, containing identifiers (perhaps names or registration numbers) for all authorized voters, along with the voters' public keys. Each voter is assumed to have two keys, a *registration* key and a *designation* key, whose uses are described below.

Third, the tabulation tellers collectively generate a public key for a distributed encryption scheme and post it on the bulletin board. Decryption of messages encrypted under this key requires the participation of all tabulation tellers.

Finally, the registration tellers generate *credentials*, which are used to authenticate votes anonymously. Each credential is associated with a single voter. Like keys in an asymmetric cryptosystem, credentials are pairs of a public value and a private value. All public credentials are posted on the bulletin board, and each registration teller stores a share of each private credential. Private credentials can be forged or leaked only if all registration tellers collude.

3.3. Voting phase

Voters register to acquire their private credentials. Each registration teller authenticates a voter using the voter's registration key. The teller and voter then run a protocol, using the voter's designation key, that releases the teller's share of the voter's private credential to the voter. The voter combines all of these shares to construct a private credential.

Voting may take place immediately, or a long time after registration. To vote, the voter submits a private credential and a *choice* of a candidate (both encrypted), along with a proof that the vote is well-formed, to some or all of the ballot boxes. (This submission does not require either of the voter's keys.) Replication of the vote across the ballot boxes is used to guarantee availability of the vote for tabulation.

Resisting coercion. The key idea (due to Juels et al. [45]) that enables voters to resist coercion, and defeats vote selling, is that voters can substitute *fake* credentials for their real credentials, then behave however the adversary demands. For example:

If the adversary demands that the voter...	Then the voter...
Submits a particular vote	Does so with a fake credential.
Sells or surrenders a credential	Supplies a fake credential.
Abstains	Supplies a fake credential to the adversary and votes with a real one.

To construct a fake credential, the voter locally runs an algorithm to produce fake private credential shares that, to an adversary, are indistinguishable from real shares. The faking algorithm requires the voter's private designation key. The voter combines these shares to produce a fake private credential; the voter's public credential remains unchanged.

Revoting. Voters might submit more than one vote per credential. The supervisor has the flexibility to specify a policy on how to tally such revotes. If revotes are not allowed, then all votes submitted under duplicate credentials are eliminated. If revotes are allowed, then the voter must include a proof in later votes to indicate which earlier votes are being replaced. This proof must demonstrate knowledge of the credential and choice used in both votes, preventing an adversary from revoting on behalf of a voter.

Ballot design. Civitas is compatible with the use of any ballot design for which a proof of well-formedness is possible. Our prototype supports the use of ballots in which voters may choose a single candidate (plurality voting), any subset of candidates (approval voting), or a ranking of the candidates (ranked voting). However, ranked voting introduces covert channels that enable attacks on coercion resistance. We discuss this vulnerability, and how to eliminate it, in the accompanying technical report [18].¹¹

Write-in votes could also be supported by Civitas, since any write-in could be considered well-formed. However, write-ins also enable attacks on coercion resistance.¹² To our knowledge, it is not possible to eliminate this vulnerability, so we chose not to implement write-ins in our prototype.

3.4. Tabulation phase

The tabulation tellers collectively tally the election:

1. **Retrieve data.** All tabulation tellers retrieve the votes from each ballot box and the public credentials from the bulletin board.
2. **Verify proofs.** The tellers check each vote to verify the proof of well-formedness. Any vote with an invalid proof is discarded. (For efficiency, our implementation actually merges this with the next step.)
3. **Eliminate duplicates.** At most one vote is retained for each credential. Votes with duplicate credentials are eliminated according to the revoting policy.
4. **Anonymize.** Both the list of submitted votes and the list of authorized credentials are anonymized by applying a random permutation, implemented with a *mix network* [11]. In the mix, each tabulation teller in turn applies its own random permutation.
5. **Eliminate unauthorized votes.** The credentials in the anonymized votes are compared against the anon-

¹¹Other kinds of ballots can be encoded into one of these supported forms. For example, conditional ballots, in which a voter selects "yes" or "no" on some issue, then is offered particular candidates based on this selection, can be encoded as a plurality vote on a pair of a selection and a candidate.

¹²For example, the adversary could issue each voter a unique, large number, then demand that the voter submit that number as the voter's choice. If that number does not appear in the final list of decrypted choices, the adversary knows that the voter did not comply.

ymized authorized credentials. Any votes with invalid credentials are discarded.

6. **Decrypt.** The remaining choices, but not credentials, are decrypted. The final tally is publicly computable.

Verifying an election. Tabulation is made publicly verifiable by requiring each tabulation teller to post proofs that it is honestly following the protocols. All tabulation tellers verify these proofs as tabulation proceeds. An honest teller refuses to continue when it discovers an invalid proof. Anyone can verify these proofs during and after tabulation, yielding universal verifiability. A voter can also verify that his vote is present in the set retrieved by the tabulation tellers, yielding voter verifiability.

4. Security Evaluation

The Civitas voting scheme requires certain assumptions about the trustworthiness of agents and system components. We discuss what attacks are possible when these trust assumptions are violated, and what defenses an implementation of the scheme could employ.

Trust Assumption 1. *The adversary cannot simulate a voter during registration.*

There must be some period of time during which the adversary cannot simulate the voter. Otherwise the system could never distinguish the adversary from the voter, so the adversary could register and vote on behalf of a voter. Registration is a good time for this assumption because it requires authentication and can be done far in advance of the election.

During registration, Civitas authenticates voters with their registration keys. So this assumption restricts the adversary from acquiring a voter's key before the voter has registered. However, voters might attempt to sell their private registration keys, or an adversary might coerce a voter into revealing the voter's key.¹³ Both attacks violate Trust Assumption 1 by allowing the adversary to simulate a voter.

One possible defense would be to store private keys on tamper-resistant hardware, which could enforce digital non-transferability of the keys. This is not a completely effective defense, as voters could physically transfer the hardware to the adversary. Preventing such physical transfers is not generally possible, but they could be discouraged by introducing economic disincentives for voters who relinquish their keys. For example, the Estonian ID card, which contains private keys and is used for electronic voting, can be used to produce legally binding cryptographic signatures [62]. Voters would be unlikely to sell such cards, although coercion would remain a problem.

Another possible defense is to change authentication to use in-person registration as an alternative to private keys.

¹³Note that these attacks are relevant only to registration, not voting, because the voter's registration key is not used during the voting protocol.

Each registration teller would either be an *online* teller, meaning voters register with that teller remotely, or an *offline* teller, meaning voters must register in person with that teller. Offline registration tellers would be trusted to authenticate voters correctly, preventing the adversary from masquerading as the voter. At least one offline registration teller would need to exist in any election, ensuring that voters register in person with at least one teller.

For deployments of Civitas in which this trust assumption does not hold, we recommend requiring in-person registration. This compromises of our goal of a fully remote system. But it is a practical defense, since voting could still be done remotely, registration could be done far in advance of the actual election, and a single credential could be reused for multiple elections.¹⁴

Trust Assumption 2. *Each voter trusts at least one registration teller, and the channel from the voter to the voter's trusted registration teller is untappable.*

Constructing a fake credential requires the voter to modify at least one of the shares received during registration. Suppose the adversary can tap all channels to registration tellers and record the encrypted traffic between the voter and the registration tellers. Further suppose that the adversary can corrupt the voter's client so that it records all credential shares received from tellers. Then the adversary can ask the client to reveal the plaintext credential shares corresponding to the encrypted network messages. In this scenario, the voter cannot lie to the adversary about his credential shares, meaning that the voter could now sell his credential and is no longer protected from coercion. So an untappable channel is required for distribution of at least one share. The voter must also trust the teller who issued that share not to reveal it.¹⁵

An untappable channel is the weakest known assumption for a coercion-resistant voting scheme [4, 19, 37, 45, 66]. Replacing this with a more practical assumption has been an open problem for at least a decade [20]. Offline registration tellers, discussed with Trust Assumption 1, could ensure an untappable channel by supervising the registration process. Our prototype of the client employs enforced erasure of all credential shares once the voter's credential is constructed, preventing the voter from reporting shares to the adversary.

Trust Assumption 3. *Voters trust their voting clients.*

Voters enter votes directly into their clients. No mechanism ensures that the client will preserve the integrity or the confidentiality of votes. A corrupt voting client could violate coercion resistance by sending the plaintext of the voter's credential and choice to the adversary. A corrupt client could

¹⁴Such reuse would require strengthening Trust Assumptions 2 and 6 to honesty of tellers across multiple elections.

¹⁵Note that a voter must know which registration teller he is trusting, which is stronger than Trust Assumptions 5 and 6.

also violate verifiability by modifying the voter’s credential or choice before encrypting it.

Clients could be corrupted in many ways. The machine, including the network connection, could be controlled by the adversary. Any level of the software stack, from the operating system to the client application, could contain vulnerabilities or be corrupted by malicious code. The adversary might even be an insider, compromising clients during their development and distribution.

Current research aims to solve this problem by changing how voters enter their votes [12, 43, 49, 75]. The voting client is decomposed into multiple (hardware and software) components, and the voter interacts with each component to complete the voting process. For example, voting might require interacting with a smart card to obtain a randomized ballot, then interacting with a client to submit a vote on that ballot.¹⁶ Now the voter need not trust a single client, but instead that the components implementing the client will not collude. Complementary research aims to leverage trusted computing technology [72]. For example, attestation could be used to prove that no level of the hardware or software stack has been changed from a trusted, pre-certified configuration. Integrating these kinds of defenses into Civitas is important future work.

Note that this trust assumption does not require all voters to trust a single client implementation. Rather, voters may choose which client they trust. This client could be obtained from an organization the voter trusts, such as their own political party or another social organization. These organizations are free to implement their own Civitas client software on their own hardware, and to make their source code publicly available. This freedom improves upon current direct-recording electronic (DRE) voting systems, in which voters are often forced by local election authorities to use particular proprietary (or closed-source) clients that are known to contain vulnerabilities [46, 48, 73]. Another advantage over DREs is that diverse clients, provided by several organizations, could reduce the incentive to attack Civitas by raising the cost of mounting an attack.

Requiring trusted voter clients compromises our goal of a remote voting system. Even if voters download a client from a trusted organization, the software stack on a voter’s machine might not be trustworthy. Thus voters might need to travel to a location where an organization they trust has provided a client application running on a trustworthy hardware and software platform.

Trust Assumption 4. *The channels on which voters cast their votes are anonymous.*

Without this assumption, the adversary could observe network traffic and learn which voters have voted, trivially

¹⁶Another example is the use of paper as one of the components. However, this is incompatible with remote electronic voting.

violating coercion resistance—although the adversary still could not learn the voter’s choice or credential.

Our prototype of Civitas does not implement its own anonymous channel because the construction of trustworthy anonymous channels is an orthogonal research problem. It seems likely that existing anonymizing networks, such as Tor [26], would suffice if made sufficiently reliable.¹⁷

Trust Assumption 5. *At least one of the ballot boxes to which a voter submits his vote is correct.*

A correct ballot box returns all the votes that it accepted to all the tabulation tellers. This is weaker than the standard assumption (less than a third of the ballot boxes fail) made for Byzantine fault tolerance [10] and multi-party computation [33], which both require more expensive protocols.

Trust Assumption 6. *There exists at least one honest tabulation teller.*

If all the tellers were corrupted, then the adversary could trivially violate coercion resistance by decrypting credentials and votes. This assumption is not needed for verifiability, even if all the tellers collude or are corrupted—the proofs posted by tellers during tabulation will reveal any attempt to cheat. Fault tolerance techniques [14, 30] would increase the difficulty of corrupting all the tellers.

Attacks on election authorities. Trust Assumptions 2, 5, and 6 allow all but one election authority of each kind to be corrupted. But certain attacks might still be mounted:

- A corrupt registration teller might fail to issue a valid credential share to a voter. The voter can detect this, but coercion resistance requires that the voter cannot prove that a share is valid or invalid to a third party. Defending against this could involve the voter and another election authority, perhaps an external auditor, jointly attempting to re-register the voter. The auditor could then attest to the misbehavior of a registration teller.
- The bulletin board might attempt to alter messages. But this is detectable since messages are signed. A bulletin board might also delete messages. This is an attack on availability, which is addressed in Section 10.
- A corrupt registrar might add fictitious voters or remove legitimate voters from the electoral roll. Each tabulation teller can defend against this by refusing to tabulate unless the electoral roll is correct according to some external policy.
- A corrupt supervisor might post an incorrect ballot design, stop an election early, or even attempt to simulate an election with only one real voter. Voters and tabulation tellers should cease to participate in the election once the supervisor exhibits such behavior.

¹⁷A vote typically fits into just three packets, so scalability and timing attacks seem unlikely to present problems.

All election authorities might be simultaneously corrupted if they all run the same software. For example, an insider working at the software supplier might hide malicious code in the tabulation teller software. As discussed in Trust Assumption 6, this attack could violate coercion resistance, but it could not violate verifiability. To defend against insider attacks, election authorities should use diverse implementations of the Civitas protocols.

Trust Assumption 7. *The Decision Diffie-Hellman (DDH) and RSA assumptions hold, and SHA-256 implements a random oracle.*

DDH and RSA are standard cryptographic assumptions. The more fundamental assumption for Civitas is DDH, as the JCJ security proof is a reduction from it.

5. Cryptographic Components

Civitas uses many cryptographic components. This section gives an overview of these; the accompanying technical report [18] contains a detailed specification of the protocols. Many components require posting messages to the bulletin board. These messages must be signed by the poster. Also, a variety of zero-knowledge proofs are used to enforce the honest execution of protocols. These proofs are made non-interactive via the Fiat-Shamir heuristic [29], so their security is in the random oracle model [5]. Civitas implements a random oracle with SHA-256.

Security proof. The security of Civitas follows from the JCJ security proof [45] and the individual security proofs of each component, cited below. We give a security proof for the registration protocol in the accompanying technical report [18].

5.1. Setup phase

Keys. The supervisor posts RSA public keys representing the election authorities. These keys are used for authentication of agents and messages. The choice of RSA is for convenience, since many real-world organizations already have RSA keys, but could be replaced by another cryptosystem. The tabulation tellers also generate a distributed El Gamal public key, described below. The registrar posts each voter’s registration public key (RSA, again for convenience) and designation public key (El Gamal).

Encryption scheme. Civitas implements a distributed El Gamal scheme similar to Brandt’s [7]. The supervisor posts a message (p, q, g) describing the cryptosystem parameters: a prime $p = 2kq + 1$, where q is also prime, and a generator g of the order q subgroup of \mathbb{Z}_p^* . This subgroup, denoted \mathcal{M} , is the message space of the cryptosystem. The tabulation tellers generate an El Gamal public key K_{TT} for which each teller holds a share of the corresponding private key.

Encryption of message m under key K with randomness r is denoted $\text{Enc}(m; r; K)$. We omit r or K from this notation when they are unimportant or clear from context. Decryption of a ciphertext c that was encrypted under key K_{TT} , denoted $\text{Dec}(c)$, requires all tabulation tellers.

El Gamal encryption is *homomorphic* with respect to multiplication. That is, $\text{Enc}(m) \cdot \text{Enc}(n) = \text{Enc}(m \cdot n)$. El Gamal permits a probabilistic *reencryption* operation, denoted $\text{Reenc}(c)$ for a ciphertext c , which produces a new encryption of the same plaintext. Encryption can be made *non-malleable*, preventing the use of homomorphisms and reencryption, by the use of Schnorr signatures [68]. Civitas uses non-malleable encryption until the tabulation phase, where malleability is required.

Civitas uses two zero-knowledge proofs to ensure the honesty of tellers during key generation and during decryption. The first is a *proof of knowledge of a discrete logarithm* due to Schnorr [67]. The second is a *proof of equality of discrete logarithms* due to Chaum and Pedersen [13].

Credential generation. Civitas uses a novel construction for credentials, based on ideas found in earlier work [20, 37, 45]. The security of this construction is proved in the accompanying technical report [18].

For each voter, each registration teller i individually generates a random element of \mathcal{M} as private credential share s_i . The corresponding public share S_i is $\text{Enc}(s_i; K_{\text{TT}})$. The registration teller posts S_i on the bulletin board and stores s_i for release during registration. After all tellers have posted a share, the voter’s public credential S is publicly computable as $\prod_i \text{Enc}(s_i; K_{\text{TT}})$, which by the homomorphic property is equal to $\text{Enc}(\prod_i s_i; K_{\text{TT}})$.

5.2. Voting phase

Registration. To acquire a private credential, a voter contacts each registration teller. The voter authenticates using his registration key, then establishes a shared AES session key using the Needham-Schroeder-Lowe [51] protocol. The voter requests registration teller i ’s share s_i of the private credential. The registration teller responds with (s_i, r, S'_i, D) , where r is random, $S'_i = \text{Enc}(s_i; r; K_{\text{TT}})$ and D is a *designated-verifier reencryption proof* (DVRP) due to Hirt and Sako [37]. The proof shows that S'_i is a reencryption of S_i , the public credential share. Construction of this proof requires the voter’s public designation key. The voter verifies that S'_i was computed correctly from s_i and r , then verifies the DVRP. These verifications convince the voter, and only the voter, that the private share is correct with respect to the public share posted on the bulletin board—i.e., that S_i is an encryption of s_i . After retrieving all the shares, the voter constructs private credential s , where $s = \prod_i s_i$.

Voting. To cast a vote, a voter posts an unsigned message $\langle \text{Enc}(s; K_{\text{TT}}), \text{Enc}(v; K_{\text{TT}}), P_w, P_k \rangle$ to some or all of

the ballot boxes, where s is the voter’s private credential, v is the voter’s choice, and P_w and P_k are zero-knowledge proofs. P_w , implemented with a *1-out-of- L reencryption proof* due to Hirt and Sako [37], shows that the vote is well-formed with respect to the ballot design of the election. Given $C = \{c_i \mid 1 \leq i \leq L\}$ and c , this reencryption proof shows there exists an i such that $c_i = \text{Reenc}(c)$. P_k , implemented by adapting a proof due to Camenisch and Stadler [9], shows that the submitter simultaneously knows s and v . This defends against an adversary who attempts to post functions of previously cast votes.

Resisting coercion. To construct a fake credential, a voter chooses at least one registration teller and substitutes a random group element $s'_i \in \mathcal{M}$ for the share s_i that registration teller sent to the voter. The voter can construct a DVRP that causes this fake share to appear real to the adversary, unless the adversary has corrupted the registration teller the voter chose (in which case the adversary already knows the real share), or unless the adversary observed the channel used by the registration teller and voter during registration (in which case the adversary has seen the real proof). By Trust Assumption 2, there exist some teller and channel that the adversary does not control, so it is always possible for voters to fake credentials.

5.3. Tabulation phase

Ballot boxes. Recall from Section 3 that ballot boxes are instances of an insert-only log service. Ballot boxes have one additional function, reporting their contents at the end of an election. When the supervisor closes the election, each ballot box posts a commitment to its contents on the bulletin board. The supervisor then posts his own signature on all these commitments, defining the set of votes to be tabulated. Thus, if a voter posts a vote to at least one correct ballot box, the vote will be tabulated.¹⁸ Note that ballot boxes do not check validity of votes.

Since ballot boxes operate independently, never contacting other ballot boxes, this ballot box construction scales easily. Moreover, this construction ensures that all votes are available for tabulation—a requirement of universal verifiability—without expensive fault tolerance protocols.

Mix network. A mix network is used to anonymize submitted votes and authorized credentials. Civitas implements a reencryption mix network made verifiable by randomized partial checking [40], in which each teller in the network performs two permutations.¹⁹

¹⁸A malicious supervisor could violate this by excluding a correct ballot box. This trust in the supervisor could be eliminated by using a more expensive agreement protocol.

¹⁹Randomized partial checking reveals some small amount of information about these permutations. In the worst case, when all but one teller is corrupted, the size of the set within which a vote or credential is anonymous

Duplicate and invalid credential elimination. It would be easy to eliminate votes containing duplicate or invalid credentials if credentials could be decrypted. However, this would fail to be coercion-resistant, because voters’ private credentials would be revealed. Instead, a zero-knowledge protocol called a *plaintext equivalence test* (PET) is used to compare ciphertexts. Given c and c' , a PET reveals whether $\text{Dec}(c) = \text{Dec}(c')$, but nothing more about the plaintexts of c and c' . Civitas implements a PET protocol due to Jakobsson and Juels [39]. For duplicate elimination, a PET must be performed on each pair of submitted credentials. Similarly, to eliminate invalid credentials, PETs must be performed to compare each submitted credential with every authorized credential.²⁰ These pairwise tests cause credential elimination to take quadratic time.

6. Scalability

There are two main challenges for scalability in Civitas. First, elimination of duplicate and invalid credentials takes quadratic time. Second, tabulation requires each teller to perform computation for each vote.

Our solution to both challenges is to group voters into *blocks*, which are virtual precincts. Like real-world precincts, the tally for each block can be computed independently, block results are public, and voters are anonymous within their block. Unlike real-world precincts, the assignment into blocks need not be based on physical location. For example, voters might be assigned to blocks in a way that is verifiably pseudorandom, reducing the risk of reprisal by the adversary against an entire block of voters. Blocking also enables the production of *early returns*, in which a fraction of blocks are tabulated to predict the outcome of the election.

Implementing blocking is straightforward. The registrar publicly assigns each voter to a block. Each submitted vote identifies, in plaintext, the block in which its credential (supposedly) resides. Vote proof P_k is extended to make this identifier non-malleable.

Without blocking, duplicate elimination requires $O(N^2)$ PETs, where N is the number of all submitted votes. With blocking, $O(BM^2)$ PETs are required, where $B = \lfloor \frac{V}{K} \rfloor$ is the number of blocks, V is the number of voters, K is the minimum number of voters per block, and M is the maximum number of votes submitted in a block. Likewise, blocking reduces invalid credential elimination from

is halved. By a result of Gomułkiewicz et al. [34], the revealed information can be made statistically small by requiring each teller to perform a total of five permutations. We estimate this would increase tabulation time by at most 3%. Mix networks based on zero-knowledge proofs [32, 57] would improve anonymity at the cost of more expensive verification.

²⁰The presence of invalid credentials is an information channel. For example, if there are zero invalid credentials, then no voter submitted a vote with a fake credential. The adversary could detect this from the PET results posted on the bulletin board. To eliminate this channel, each teller could post a random number of votes with invalid credentials.

Table 1. Modular exponentiations per block

Agent	Action	Protocol	BB
RT	Generate all credentials	$4K$	K
	Distribute all credentials	$14K$	-
Voter	Retrieve a credential	$12A$	A
	Vote	$4C + 7$	-
TT	Retrieve data	-	$AK + A + 1$
	Verify proofs	$4M(C + 1)$	-
	Eliminate duplicates	$\binom{M}{2}(8A - 1)$	$3A$
	Anonymize (mixes)	$2(A + 1)(M + K)$	$2A$
	Eliminate invalids	$KM(8A - 1)$	$3A$
	Decrypt	$K(4A - 1)$	A

$O(VN)$ PETs to $O(BKM)$. The B factor in each of these terms is easily parallelizable, since a different set of machines can be used to implement the tabulation tellers for each block. Tabulation time then depends on M and K , but not V . Therefore performance can scale independently of the number of voters.

Table 1 identifies the number of modular exponentiations performed per block by individual agents: registration tellers (RT), tabulation tellers (TT), and voters. (Tabulation time is dominated by modular exponentiations.) The table distinguishes *protocol* exponentiations, which are required by the Civitas voting scheme regardless of the implementation of the bulletin board, from *bulletin board* (BB) exponentiations, which are required by the particular implementation used in our prototype. BB exponentiations result from RSA signatures and verifications. Exponentiations are counted under the assumption that there are no duplicate votes and that no voters abstain, maximizing the number of PETs required. Parameter A describes the number of election authorities of each kind—i.e., if $A = 4$, then there are four registration tellers, four tabulation tellers, and four ballot boxes. Regardless of A , there is a single bulletin board. Table 1 assumes a plurality ballot with C candidates.

7. Implementation in Jif

Our prototype of Civitas is implemented in Jif_E [15], an extension of Jif 3.0 [54, 56]. Jif is a security-typed language in which programs are annotated with information-flow security policies. The Jif compiler and runtime guarantee end-to-end enforcement of these policies. Information-flow policies control both the release and propagation of information, enabling the protection of both sensitive data and data derived therefrom. Information-flow policies are therefore stronger than access control policies, which control only the release of information.

Jif security policies are expressed using the *decentralized label model* [55], which allows specification of confidentiality and integrity requirements of principals. Such policies

are useful for constructing systems like Civitas, in which principals need to cooperate yet are mutually distrusting. For example, if information is labeled with confidentiality policy $RT_1 \rightarrow voter_{76}$, then principal RT_1 permits principal $voter_{76}$ to learn the information; such a policy would be suitable for the private credential share generated by registration teller RT_1 for $voter_{76}$. Similarly, if information is labeled with integrity policy $TT_3 \leftarrow Sup$, then principal TT_3 requires that only principal Sup has influenced the information; such a policy would be suitable for the ballot design, which only the supervisor may specify.

In general, a principal p may specify a set R of readers in confidentiality policy $p \rightarrow R$. Jif_E extends Jif with *declassification* and *erasure* policies [16], which allow principals to state *conditions* on when the set of readers in a confidentiality policy may be changed.

Declassification policies allow the set of readers of information to be expanded. For example, in the implementation of mix networks, each tabulation teller must commit to random bits. The bits are then revealed and used to verify the mix. The security of the mix requires maintaining the secrecy of these bits until all tellers have committed. In the code, this requirement is expressed using a declassification policy. The policy annotating the variable storing the random bits of TT_i indicates that the information is readable only by TT_i until condition *AllCommitted* is satisfied, upon which the information may be declassified to be readable by all principals. *AllCommitted* becomes true at the program point where all commitments have been received.

Erasure policies mandate conditions upon which the set of readers must be restricted. For example, each registration teller must store a private credential share for each voter until the voter requests it. After this, the teller may erase the share, ensuring that the share cannot later be disclosed.²¹ In

²¹Erasure is a design choice that impacts recovery from voters' accidental loss or deletion of credentials. If tellers do not erase shares, then tellers can reissue credentials. But if tellers do erase shares, then reissue is not possible. Instead, tellers would need to revoke lost shares and issue new shares. This is left as future work.

Table 2. Lines of Jif_E code per component

Component	LOC
Tabulation teller	5,740
Common	3,173
Registration teller	1,290
Supervisor	1,138
Log service (bulletin board and ballot box)	911
Voter client	826
Registrar	308
Total	13,386

the code, the variable storing the share is annotated with an erasure policy indicating that this information becomes unreadable by all principals when condition *Delivered* is satisfied. *Delivered* becomes true at the program point where receipt of the share has been acknowledged by the voter. The Jif_E compiler inserts code at that point to erase the information from memory.

Our implementation of Civitas totals about 13,000 lines of Jif_E code. Table 2 gives the number of lines of code in each component; common code includes shared data structures and utility methods for retrieving and caching election information. About 8,000 additional lines of Java code are used to perform I/O and to implement number-theoretic operations such as encryption and zero-knowledge proofs.

8. Performance

A voting system is practical only if tabulation can be completed in reasonable time, with reasonable cost and security. Civitas offers a tradeoff between these three factors, because tabulation can be completed more quickly by accepting higher cost or lower security.

Notions of reasonable time, cost, and security may differ depending on the election or the observer. In current U.S. elections, accurate predictions of election results are available within a few hours. Therefore, we chose a target tabulation time of five hours. The two most important parameters affecting security are K , the minimum number of voters within each block, and A , the number of authorities of each kind.²² As reasonable values for these parameters, we chose $K = 100$ and $A = 4$. Anonymity within 100 voters seems comparable to what is available in current real-world elections, where results are tabulated at a precinct level and observers might correlate voters with ballots.²³ Similarly, four mutually distrusting authorities might offer better oversight than real-world elections.

²²Recall from Section 6 that if $A = 4$, then there are four registration tellers, four tabulation tellers, and four ballot boxes.

²³Random block assignment might even offer stronger anonymity than real-world elections.

Experiment design. We used Emulab [74] as an experiment testbed. The experiments ran on machines containing 3.0 GHz Xeon processors and 1 GB of RAM, networked on a 1 Gb LAN. Note that only tabulation tellers actually need hardware this fast, whereas voters could use substantially less powerful hardware without impacting performance or the voting experience. Our machines ran Red Hat Linux 9.0 and Java 1.5.0_11. For RSA, AES, and SHA implementations, we used Bouncy Castle JCE provider 1.33. We implemented the remaining cryptographic functionality, including El Gamal and zero-knowledge proofs, ourselves. We used a C library, GMP 4.2.1, for implementations of modular exponentiation and multiplication.

Key lengths were chosen to meet or exceed NIST recommendations for 2011–2030 [3]. We used 128-bit AES keys, 2048-bit RSA keys, and 224-bit El Gamal keys from a 2048-bit group—i.e., $|p| = 2048$ and $|q| = 224$. A modular exponentiation in this size group required about 3.7 ms.

Each experiment simulated all phases of a complete election, including all the cryptographic protocols in Section 5. Therefore the results should be representative of a real deployment. All experiments used plurality ballots with three candidates. No voters abstained, so $N \geq V$ and $M \geq K$.²⁴ Experiments were repeated three times, and we report the sample mean. The sample standard deviation was always less than 2% of the mean.

Setup and voting time. Generation of keys and credentials scales linearly in the number of authorities and voters, respectively, and can be conducted offline. During the voting phase, voters retrieve credential shares from registration tellers and submit votes to ballot boxes. A voter client takes about 325 ms to acquire a credential share from a registration teller, and about 20 ms to submit a vote to a ballot box. Thus, for four authorities, it takes a voter less than 1.4 seconds to retrieve credentials and submit a vote. From the registration teller’s perspective, it takes about 200 ms of CPU time to distribute a single voter’s credential share. A registration teller could therefore process 18,000 voters per hour.

Tabulation time and space. Figure 2(a) shows the results of four tabulation tellers processing blocks sequentially, where V is a multiple of K . The data indicate that Civitas requires 39 seconds per voter per authority to tabulate a single block, and that votes from 500 voters, in blocks of 100, can be tabulated in five hours. (The time to combine the block tallies is negligible.) Parameters A and K have non-linear effects on tabulation time, as shown in Figure 2(b) and Figure 2(c). Communication increases quadratically in A , and PETs take time proportional to K^2 . Figure 2(c) indicates that a block of 200 voters can be tabulated in less than five hours.

²⁴Recall that N is the number of votes submitted and M is the maximum number of votes submitted in a block.

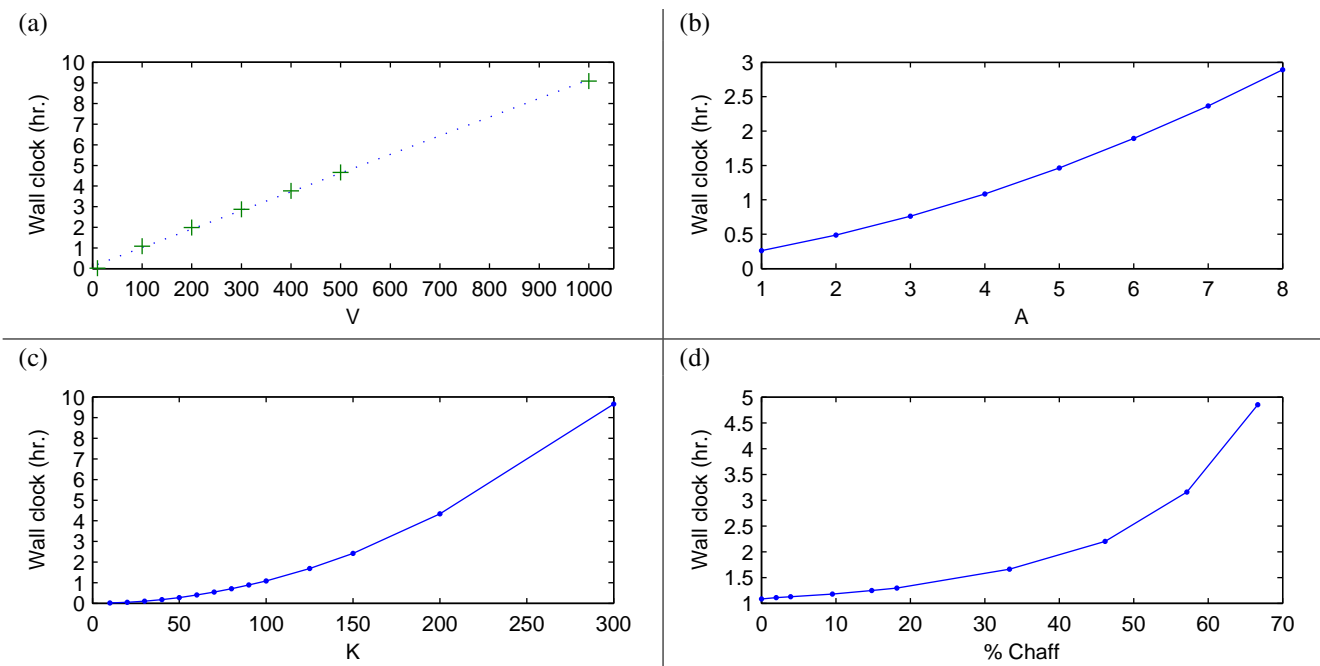


Figure 2. Tabulation time vs. (a) Voters: $K = 100, A = 4$; **(b) Authorities:** $K = V = 100$; **(c) Anonymity:** $V = K, A = 4$; **(d) Chaff:** $K = V = 100, A = 4$

The independence of blocks can be exploited to decrease tabulation time by processing blocks in parallel. Given a set of tabulation teller machines for each block, the data in Figure 2(a) predict that tabulation could be completed in about 65 minutes, independent of V . Because of the linear trade-off between time and machines at the granularity of blocks, the remaining measurements in this study are for tabulation of a single block.

The memory footprint of Civitas is very small. With $M = 100$, the active set of a tabulation teller is never more than 8 MB. The size of the active set scales linearly in M , so modern machines could easily fit tabulation in memory for substantially larger values of M (and of K , since $K \leq M$). The storage space needed for the entire bulletin board is less than 620 MB for an election where $K = 100, V = 100$, and $A = 4$. Our prototype uses a verbose XML-like message format, so we expect that storage space requirements could be reduced significantly.²⁵

Chaff. We refer to votes containing invalid and duplicate credentials as *chaff* because they are eliminated during tabulation. Because chaff increases the number of votes in a block, it increases tabulation time similarly to increasing anonymity parameter K . Figure 2(d) shows how tabulation

²⁵Note that voters do not need to download the entire bulletin board to verify inclusion of their votes. Rather, a voter would need to download only the list of votes (about 160 kB) used as input to the tabulation protocol, then check that his vote is in this list.

time varies as a function of the percentage of chaff votes in each block. With fraction c chaff (split between invalid and duplicate credentials), there are $M = \frac{V}{1-c}$ votes in a block. All the other graphs in this study assume $c = 0$.

Cost. A government election in a stable Western democracy currently costs \$1 to \$3 per voter [38]. Civitas would increase the cost of computing equipment but could reduce the costs associated with polling places and paper ballots. A dual-core version of our experiment machines is currently available for about \$1,500, so the machine cost to tabulate votes from 500 voters in five hours (with $K = 100$ and $A = 4$) is at worst \$12 per voter, and this cost could be amortized across multiple elections. Moving to multicore CPUs would also be likely to reduce tabulation time, since tabulation is CPU-bound (utilization is about 70–85% during our experiments), has a small memory footprint, and can be split into parallel threads that interact infrequently. Costs could be reduced dramatically if trust requirements permit a tabulation teller to lease compute time from a provider.²⁶ One provider currently offers a rate of \$1 per CPU per hour on processors similar in performance to our experiment ma-

²⁶Essentially, this means trusting the provider with the teller’s El Gamal private key share for that election so the provider can compute decryption shares. To avoid giving the provider the key share, computation might be split between the provider and teller, with the teller computing only these decryption shares. This would result in the teller performing only about 10% of the total number of modular exponentiations.

chines [71]. At this rate, tabulation for 500 voters would cost about 4¢ per voter—clearly in the realm of practicality.

Reducing security parameters also reduces cost. For example, halving K approximately quarters tabulation time. So for a ten-hour, $K = 50$, $A = 3$ election, the cost per voter would be about ten times smaller than a five-hour, $K = 100$, $A = 4$ election. El Gamal key lengths also have a significant impact. Figure 2(c) shows that, for 224-bit keys from a 2048-bit group, K can be as high as 200 while maintaining a tabulation time of under five hours. With 160-bit keys from a 1024-bit group (secure, according to NIST, from 2007–2010 [3]), K can be increased to 400. Using 256-bit keys from a 3072-bit group (secure until after 2030) currently requires decreasing K to 125.

Real-world estimates. In the 2004 general election for President of the United States, just under 2.3 million votes were reported by the City of New York Board of Elections [17]. Using the worst-case estimate we developed above, \$12 per voter, the one-time hardware cost for using Civitas to tabulate this election would be at most \$27.6 million. In comparison, Diebold submitted an estimate in 2006 of \$28.7 million in one-time costs to replace the city’s mechanical lever voting machines with optical scan machines [23]; hardware and software costs accounted for \$10.2 million of this estimate [24]. Although we cannot make any strong comparisons, the cost of Civitas does seem to be about the same order of magnitude.

9. Related Work

Voting schemes. Cryptographic voting schemes can be divided into three categories, based on the technique used to anonymize votes: homomorphic encryption [6, 20, 37, 65], blind signatures [31, 58, 59], and mix networks [4, 11, 52, 66]. JCJ and Civitas are both based on mix networks.

To optimize JCJ, Smith [70] proposes replacing PETs with reencryption into a deterministic, distributed cryptosystem. However, the proposed construction is insecure. The proposed encryption function is $\text{Enc}(m; z) = m^z$, where z is a secret key distributed among the tellers. But to test whether s is a real private credential, the adversary can inject a vote using s^2 as the private credential. After the proposed encryption function is applied during invalid credential elimination, the adversary can test whether any submitted credential is the square of any authorized credential. If so, then s is real with high probability. Araújo et al. [1] are studying another possible replacement for PETs, based on group signatures.

Civitas differs from JCJ in the following ways:

- JCJ assumes a single trusted registration authority; Civitas factors this into a registrar and a set of mutually distrusting registration tellers. As part of this, Civitas introduces a construction of credential shares.

- JCJ does not specify a means of distributing credentials; Civitas introduces a protocol for this and proves its security.
- JCJ has voters post votes to the bulletin board; Civitas introduces ballot boxes for vote storage.
- JCJ supports plurality voting; Civitas generalizes this to include approval and ranked voting methods.
- JCJ left many of the cryptographic components described in Section 5 unspecified (though JCJ also provided helpful suggestions for possible implementations); Civitas provides concrete instantiations of all the cryptographic components in the voting scheme.
- JCJ, as a voting scheme, did not study the scalability of tabulation or conduct experiments; Civitas, as both a scheme and a system, introduces blocking, studies its scalability, and reports experimental results.

Voting systems. To our knowledge, Civitas offers stronger coercion resistance than other implemented voting systems. Sensus [21], based on a blind signature scheme known as FOO92 [31], offers no defense against coercion. Neither does EVOX [36], also based on FOO92. Both systems allow a single malicious election authority to vote on behalf of voters who abstain. EVOX-MA [27] addresses this by distributing authority functionality. REVS [44, 50] extends EVOX-MA to tolerate failure of distributed components, but does not address coercion. ElectMe [69] is based on blind signatures and claims to be coercion resistant, but it assumes the adversary cannot corrupt election authorities. If the adversary learns the ciphertext of a voter’s “ticket,” the scheme fails to be receipt-free. ElectMe also is not universally verifiable. Voters can verify their votes are recorded correctly, but the computation of the tally is not publicly verifiable. Adder [47] implements a homomorphic scheme in which voters authenticate to a “gatekeeper.” If the adversary were to corrupt this single component, then Adder would fail to be coercion-resistant.

Kiayias [47] surveys several voting systems from the commercial world. These proprietary systems do not generally make their implementations publicly or freely available, nor do they appear to offer coercion resistance. The California top-to-bottom review [73] of commercial electronic voting systems suggests that these systems offer completely inadequate security.

The W-Voting system [49] offers limited coercion resistance. It requires voters to sign votes, which appears susceptible to attacks in which a coercer insists that the voter abstain or submit a vote prepared by the coercer. It also allows voters to submit new votes, which replace older votes. So unlike Civitas, an adversary could successfully coerce a voter by forcing the voter to submit a new vote, then keeping the voter under surveillance until the end of the election.

Prêt à Voter 2006 [64] offers a weak form of coercion resistance, if voting is supervised. The construction of ballots depends on non-uniformly distributed seeds, which might enable the adversary to learn information about how voters voted. In remote settings, Prêt à Voter offers no coercion resistance. The adversary, by observing the voter during voting, will learn what vote was cast.

VoteHere [57] offers coercion resistance, assuming a supervised voting environment. Removing this assumption seems non-trivial, since the supervised environment includes a voting device with a trusted random number generator. This generator could be subverted in a remote setting, enabling the adversary to learn the voter's vote.

The primary goal of Punchscan [61] is high integrity verification of optical scan ballots. Punchscan does not claim to provide coercion resistance. Instead, under the assumption that voting takes place in a supervised environment, Punchscan offers a weaker property: The adversary learns nothing by observing data revealed during tabulation. This assumption rules out coercion-resistant remote voting. For confidentiality, Punchscan assumes that the election authority is not corrupted, even partially, by the adversary.

10. Toward a Secure Voting System

Some open technical problems must be solved before Civitas, or a system like it, could be used to secure national elections. Two such problems are that Civitas assumes a trusted voting client, and that in practice, the best way to satisfy two of the Civitas trust assumptions is in-person registration.

We did not address availability in this work. However, the design of Civitas accommodates complementary techniques for achieving availability. To improve the availability of election authorities, they could be implemented as Byzantine fault-tolerant services [10, 60]. Also, the encryption scheme used by Civitas could be generalized from the current distributed scheme to a threshold scheme. This would enable election results to be computed even if some tabulation tellers become unresponsive or exhibit faulty behavior, such as posting invalid zero-knowledge proofs.²⁷ For a threshold scheme requiring k out of n tabulation tellers to participate in decryption, no more than $k - 1$ tellers may be corrupted, otherwise coercion resistance could be violated. For availability, a new trust assumption must be added: At least k tellers do not fail.²⁸

Application-level denial of service is particularly problematic, because an adversary could insert chaff to inflate tabulation time. A possible defense, in addition to standard techniques such as rate-limiting and puzzles, would be to

²⁷Recovery from these faults would need to ensure that the adversary cannot exploit any partial information from aborted subphases.

²⁸The adversary could increase tabulation time by forcing at most $n - k$ restarts. But as long as no more than $k - 1$ tellers are corrupted, the adversary cannot successfully cause tabulation to be aborted.

require a *block capability* in each submitted vote. The adversary would need to learn the capability for each block, individually, to successfully inflate tabulation time for that block. Another possible defense is to weaken coercion resistance so that chaff votes could be detected without requiring PETs. These defenses are left as future work.

We have not investigated the usability of Civitas, although usability is more important than security to some voters [35]. Management of credentials is an interesting problem for the use of Civitas. Voters might find generating fake credentials, storing and distinguishing real and fake credentials (especially over a long term), and lying convincingly to an adversary to be quite difficult. Recovery of lost credentials is also an open problem.

There are open non-technical problems as well; we give three examples. First, some people believe that any use of cryptography in a voting system makes the system too opaque for the general public to accept.²⁹ Second, remote electronic voting requires voters to have access to computers, but not all people have such access now. Third, some real-world attacks, such as attempts to confuse or misinform voters about the dates, significance, and procedures of elections, are not characterized by formal security models. Mitigation of such attacks is important for real-world deployments, but beyond the scope of this paper.

Finally, a report on the security of a real-world remote voting system, SERVE, identifies a number of open problems in electronic voting [42]. These problems include transparency of voter clients, vulnerability of voter clients to malware, and vulnerability of the ballot boxes to denial-of-service attacks that could lead to large-scale or selective disenfranchisement. However, Civitas does address other problems raised by the report: the voter client is not a DRE, trust is distributed over a set of election authorities, voters can verify their votes are counted, spoofing of election authorities is not possible due to the use of digital signatures, vote buying is eliminated by coercion resistance, and election integrity is ensured by verifiability.

11. Conclusion

This paper describes the design, implementation, and evaluation of Civitas, a remote voting system whose underlying voting scheme is proved secure under carefully articulated trust assumptions. To our knowledge, this has not been done before. Civitas provides stronger security than previously implemented electronic voting systems. Experimental results show that cost, tabulation time, and security can be practical for real-world elections.

Civitas is based on a previously-known voting scheme, but elaborating the scheme into an implemented system led

²⁹Our stance is that it is unnecessary to convince the general public directly. Rather, we need to convince experts by using principled techniques that put security on firm mathematical foundations.

to new technical advances: a secure registration protocol and a scalable vote storage system. Civitas thus contributes to both the theory and practice of electronic voting. But perhaps the most important contribution of this work is evidence that secure electronic voting could be made possible. We are optimistic about the future of electronic voting systems constructed, like Civitas, using principled techniques.

Website

The accompanying technical report and prototype source code are available from:

<http://www.cs.cornell.edu/projects/civitas>

Acknowledgments

We thank Michael George, Anil Nerode, Nathaniel Nystrom, Tom Roeder, Peter Ryan, Fred B. Schneider, Tyler Steele, Hakim Weatherspoon, Lantian Zheng, and Lidong Zhou for discussions about this work; Jed Liu, Tudor Marian, and Tom Roeder for consultation on performance experiments; and the anonymous reviewers for their comments. We thank the participants of FEE'05 and Frontiers of Electronic Voting (Dagstuhl Seminar 07311) for feedback on preliminary versions of this work.

References

- [1] R. Araujo, S. Foulle, and J. Traoré. On coercion-resistant schemes with linear work. In *Proc. of Frontiers of Electronic Voting: Dagstuhl Seminar 07311*, July 2007.
- [2] J. Bannet, D. W. Price, A. Rudys, J. Singer, and D. S. Wallach. Hack-a-vote: Security issues with electronic voting systems. *IEEE Security & Privacy*, 2(1):32–37, Jan. 2004.
- [3] E. Barker, W. Barker, W. Burr, W. Polk, and M. Smid. Recommendation for key management. NIST Special Publication 800-57 Part 1, Mar. 2007.
- [4] O. Baudron, P.-A. Fouque, D. Pointcheval, G. Poupard, and J. Stern. Practical multi-candidate election system. In *Proc. of ACM Symposium on Principles of Distributed Computing*, pages 274–283, Aug. 2001.
- [5] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proc. of ACM Conference on Computer and Communications Security*, pages 62–73. ACM, Nov. 1993.
- [6] J. D. C. Benaloh. *Verifiable Secret-Ballot Elections*. PhD thesis, Yale University, Sept. 1987.
- [7] F. Brandt. Efficient cryptographic protocol design based on distributed El Gamal encryption. In *Proc. of International Conference on Information Security and Cryptology*, pages 32–47, Dec. 2005.
- [8] Brennan Center for Justice. The machinery of democracy: Voting system security, accessibility, usability, and cost. New York University, Oct. 2006.
- [9] J. Camenisch and M. Stadler. Efficient group signature schemes for large groups. In *Proc. of International Cryptology Conference (CRYPTO)*, pages 410–424, Aug. 1997.
- [10] M. Castro and B. Liskov. Practical Byzantine fault tolerance. In *Proc. of Symposium on Operating System Design and Implementation*, pages 173–186, Feb. 1999.
- [11] D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, 1981.
- [12] D. Chaum. SureVote. <http://www.surevote.com>, 2007. International patent WO 01/55940 A1, 02 August 2001.
- [13] D. Chaum and T. P. Pedersen. Wallet databases with observers. In *Proc. of International Cryptology Conference (CRYPTO)*, pages 89–105, Aug. 1992.
- [14] L. Chen and A. Avizienis. N-version programming: A fault tolerance approach to reliability of software operation. In *International Symposium on Fault-Tolerant Computing*, 1978.
- [15] S. Chong and A. C. Myers. End-to-end enforcement of erasure. In submission.
- [16] S. Chong and A. C. Myers. Language-based information erasure. In *Proc. of IEEE Computer Security Foundations Workshop*, pages 241–254, June 2005.
- [17] City of New York Board of Elections. General election results. <http://www.vote.nyc.ny.us/results.html>, 2004.
- [18] M. R. Clarkson, S. Chong, and A. C. Myers. Civitas: Toward a secure remote voting system. Technical Report 2007-2081, Cornell University, May 2007. Revised Mar. 2008. <http://hdl.handle.net/1813/7875>.
- [19] R. Cramer, M. Franklin, B. Schoenmakers, and M. Yung. Multi-authority secret-ballot elections with linear work. In *Proc. of International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, pages 72–83, May 1996.
- [20] R. Cramer, R. Gennaro, and B. Schoenmakers. A secure and optimally efficient multi-authority election scheme. In *Proc. of International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, pages 103–118, May 1997.
- [21] L. F. Cranor and R. K. Cytron. Sensus: A security-conscious electronic polling system for the Internet. In *Proc. of IEEE Hawaii International Conference on Systems Science*, pages 561–570, Jan. 1997.
- [22] S. Delaune, S. Kremer, and M. Ryan. Coercion-resistance and receipt-freeness in electronic voting. In *Proc. of IEEE Computer Security Foundations Workshop*, pages 28–42, July 2006.
- [23] Diebold Election Systems. New York City BOE voting system, Cost response: Cost proposal summary, October 20, 2006. <http://www.vote.nyc.ny.us/rfi.html>.
- [24] Diebold Election Systems. New York City BOE voting system, Cost response: Lever replacement solution: Optical scan pollsite system, October 20, 2006. <http://www.vote.nyc.ny.us/rfi.html>.
- [25] D. L. Dill, B. Schneier, and B. Simons. Voting and technology: Who gets to count your vote? *Communications of the ACM*, 46(8):29–31, Aug. 2003.
- [26] R. Dingledine, N. Mathewson, and P. F. Syverson. Tor: The second-generation onion router. In *Proc. of USENIX Security Symposium*, pages 303–320, Aug. 2004.
- [27] B. W. DuRette. Multiple administrators for electronic voting. Bachelor's Thesis, Massachusetts Institute of Technology, 1999.
- [28] D. Evans and N. Paul. Election security: Perception and reality. *IEEE Security & Privacy*, 2(1):24–31, Jan. 2004.
- [29] A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Proc. of International Cryptology Conference (CRYPTO)*, pages 186–194, Aug. 1986.
- [30] S. Forrest, A. Somayaji, and D. Ackley. Building diverse computer systems. In *Proc. of IEEE Workshop on Hot Topics in Operating Systems*, May 1997.
- [31] A. Fujioka, T. Okamoto, and K. Ohta. A practical secret voting scheme for large scale elections. In *Proc. of International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, pages 244–251, May 1992.
- [32] J. Furukawa. Efficient and verifiable shuffling and shuffle-decryption. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, E88-A(1):172–188, 2005.
- [33] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In *Proc. of ACM Symposium on Theory of Computing*, pages 218–229, 1987.
- [34] M. Gomułkiewicz, M. Klonowski, and M. Kutylowski. Rapid mixing and security of Chaum's visual electronic voting. In *Proc. of European Symposium on Research in Computer Security*, pages 132–145, 2003.

- [35] P. S. Herrnson, R. G. Niemi, M. J. Hanmer, B. B. Bederson, and F. C. Conrad. *Voting Technology: The Not-So-Simple Act of Casting a Ballot*. Brookings Institution Press, 2008.
- [36] M. Herschberg. Secure electronic voting over the world wide web. Master's thesis, Massachusetts Institute of Technology, 1997.
- [37] M. Hirt and K. Sako. Efficient receipt-free voting based on homomorphic encryption. In *Proc. of International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, pages 539–556, May 2000.
- [38] International Foundation for Election Systems. Getting to the CORE—A global survey on the cost of registration and elections, June 2006. <http://www.undp.org/governance/docs/Elections-Pub-Core.pdf>.
- [39] M. Jakobsson, A. Juels, and R. L. Rivest. Mix and match: Secure function evaluation via ciphertxts. In *Proc. of International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT)*, pages 162–177, Dec. 2000.
- [40] M. Jakobsson, A. Juels, and R. L. Rivest. Making mix nets robust for electronic voting by randomized partial checking. In *Proc. of USENIX Security Symposium*, pages 339–353, Aug. 2002.
- [41] D. Jefferson, A. D. Rubin, B. Simons, and D. Wagner. Analyzing Internet voting security. *Communications of the ACM*, 47(10):59–64, Oct. 2004.
- [42] D. Jefferson, A. D. Rubin, B. Simons, and D. Wagner. A security analysis of the secure electronic registration and voting experiment (SERVE). <http://www.servesecurityreport.org/paper.pdf>, Jan. 2004.
- [43] R. Joaquim and C. Ribeiro. CodeVoting: Protecting against malicious vote manipulation at the voter's PC. In *Proc. of Frontiers of Electronic Voting: Dagstuhl Seminar 07311*, July 2007.
- [44] R. Joaquim, A. Zúquete, and P. Ferreira. REVS—A robust electronic voting system. In *Proc. of IADIS International Conference on e-Society*, June 2003.
- [45] A. Juels, D. Catalano, and M. Jakobsson. Coercion-resistant electronic elections. In *Proc. of Workshop on Privacy in the Electronic Society*, pages 61–70, Nov. 2005.
- [46] C. Karlof, N. Sastry, and D. Wagner. Cryptographic voting protocols: A systems perspective. In *Proc. of USENIX Security Symposium*, 2005.
- [47] A. Kiayias, M. Korman, and D. Walluck. An Internet voting system supporting user privacy. In *Proc. of Annual Computer Security Applications Conference*, pages 165–174. IEEE Computer Society, Dec. 2006.
- [48] T. Kohno, A. Stubblefield, A. D. Rubin, and D. S. Wallach. Analysis of an electronic voting system. In *Proc. of IEEE Symposium on Security and Privacy*, pages 27–42, May 2004.
- [49] M. Kutylowski, F. Zagórski, et al. W-Voting system. <http://w-vote.im.pwr.wroc.pl>, 2007.
- [50] R. Lebre, R. Joaquim, A. Zúquete, and P. Ferreira. Internet voting: Improving resistance to malicious servers in REVS. In *Proc. of IADIS International Conference on Applied Computing*, Mar. 2004.
- [51] G. Lowe. An attack on the Needham-Schroeder public key authentication protocol. *Information Processing Letters*, 56(3):131–136, Nov. 1995.
- [52] E. Magkos, M. Burmester, and V. Chrissikopoulos. Receipt-freeness in large-scale elections without untappable channels. In *Proc. of IFIP Conference on E-Commerce, E-Business, E-Government*, pages 683–694, Oct. 2001.
- [53] R. Mercuri. Statement on electronic voting. <http://www.notablessoftware.com/RMstatement.html>, 2007.
- [54] A. C. Myers. JFlow: Practical mostly-static information flow control. In *Proc. of Symposium on Principles of Programming Languages*, pages 228–241, Jan. 1999.
- [55] A. C. Myers and B. Liskov. Protecting privacy using the decentralized label model. *ACM Transactions on Software Engineering and Methodology*, 9(4):410–442, Oct. 2000.
- [56] A. C. Myers, L. Zheng, S. Zdancewic, S. Chong, and N. Nystrom. Jif: Java information flow (software release). <http://www.cs.cornell.edu/jif>, July 2001.
- [57] C. A. Neff. Verifiable mixing (shuffling) of ElGamal pairs. <http://www.votehere.org/vhti/documentation/egshuf-2.0.3638.pdf>, Apr. 2004.
- [58] M. Ohkubo, F. Miura, M. Abe, A. Fujioka, and T. Okamoto. An improvement on a practical secret voting scheme. In *Proc. of Information Security Workshop*, pages 225–234, Nov. 1999.
- [59] T. Okamoto. Receipt-free electronic voting schemes for large scale elections. In *Proc. of Security Protocols Workshop*, pages 25–35, Apr. 1997.
- [60] R. A. Peters. A secure bulletin board. Master's thesis, Technische Universiteit Eindhoven, June 2005.
- [61] S. Popoveniuc and B. Hosp. An introduction to Punchscan. In *Proc. of Workshop on Trustworthy Elections*, June 2006.
- [62] Republic of Estonia. Digital signatures act (digitaalalkirja seadus). <http://www.riigiteataja.ee/ert/act.jsp?id=694375>, 2000.
- [63] A. D. Rubin. Security considerations for remote electronic voting. *Communications of the ACM*, 45(12):39–44, Dec. 2002.
- [64] P. Y. A. Ryan and S. A. Schneider. Prêt à Voter with re-encryption mixes. In *Proc. of European Symposium on Research in Computer Security*, Sept. 2006.
- [65] K. Sako and J. Kilian. Secure voting using partially compatible homomorphisms. In *Proc. of International Cryptology Conference (CRYPTO)*, pages 411–424, Aug. 1994.
- [66] K. Sako and J. Kilian. Receipt-free mix-type voting scheme—A practical solution to the implementation of a voting booth. In *Proc. of International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, pages 393–403, May 1995.
- [67] C.-P. Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991.
- [68] C.-P. Schnorr and M. Jakobsson. Security of signed ElGamal encryption. In *Proc. of International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT)*, pages 73–89, Dec. 2000.
- [69] A. M. Shubina and S. W. Smith. Design and prototype of a coercion-resistant, voter verifiable electronic voting system. In *Proc. of Conference on Privacy, Security and Trust*, pages 29–39, Oct. 2004.
- [70] W. D. Smith. New cryptographic election protocol with best-known theoretical properties. In *Proc. of Workshop on Frontiers in Electronic Elections*, Sept. 2005.
- [71] Sun Microsystems. Sun grid compute utility: Users guide. <http://www.sun.com/service/sungrid/SunGridUG.pdf>, Mar. 2007.
- [72] M. Volkamer, A. Alkassar, A.-R. Sadeghi, and S. Schulz. Enabling the application of open systems like PCs for online voting. In *Proc. of Workshop on Frontiers in Electronic Elections*, 2006.
- [73] D. Wagner and M. Bishop. Voting systems top-to-bottom review. http://www.sos.ca.gov/elections/elections_vsr.htm, 2007.
- [74] B. White, J. Lepreau, L. Stoller, R. Ricci, S. Guruprasad, M. Newbold, M. Hibler, C. Barb, and A. Joglekar. An integrated experimental environment for distributed systems and networks. In *Proc. of Symposium on Operating System Design and Implementation*, pages 255–270, Dec. 2002.
- [75] A. Zúquete, C. Costa, and M. Romão. An intrusion-tolerant e-voting client system. In *Proc. of Workshop on Recent Advances in Intrusion-Tolerant Systems*, pages 23–27, Mar. 2007.