

CIS 551 / TCOM 401

# Computer and Network Security

Spring 2008

Lecture 22

# Announcements

---

- Project 4 is Due Friday May 2nd at 11:59 PM
- Final exam:
  - Friday, May 12th. Noon - 2:00pm DRLB A6
- Topic for today:
  - Civitas: Toward a Secure Voting System  
Michael Clarkson, Stephen Chong, Andrew Myers (2008)
- Some content adapted from slides by: Michael Clarkson and Andrew Myers

# Grade Summary So Far

---

(Weighted) Average: 70%

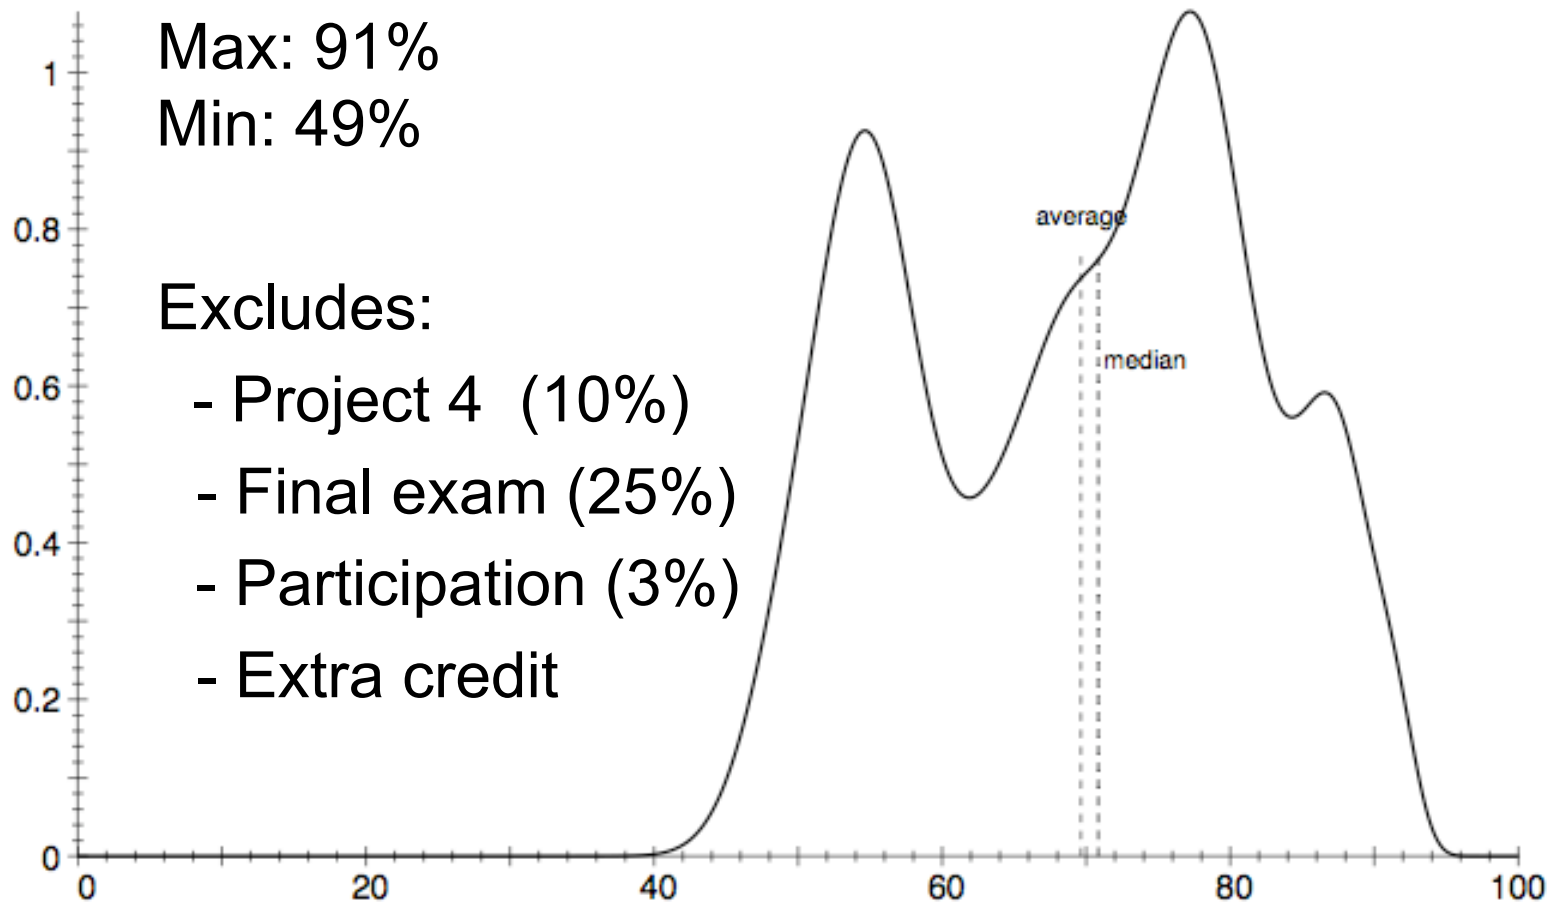
Std. Dev: 12%

Max: 91%

Min: 49%

Excludes:

- Project 4 (10%)
- Final exam (25%)
- Participation (3%)
- Extra credit



# Civitas: Secure Remote Voting

---

- A secure, remote voting system implemented at Cornell by Michael Clarkson, Steve Chong, and Andrew Myers
  - <http://www.cs.cornell.edu/projects/civitas>
- Based on an earlier voting scheme proposed by
  - Ari Juels, Dario Catalano, and Markus Jakobsson (WPES 2005)
- Terminology:
  - *Voting system*: (software) implementation
  - *Voting scheme*: cryptographic construction
  - *Voting method*: algorithm for choosing between candidates

# Remote Voting

---

- *Remote*: no supervision of voting or voters
  - Generalizes Internet voting
- The authors argue that this is the right problem to solve:
  - Does not assume supervision
  - Internet voting is common (Debian, ACM, SERVE)
  - Absentee ballots
  - Some states moving entirely to remote voting (Oregon, Washington)

# Integrity

---

- The final tally cannot be changed to be different from a public counting of the votes

*Verifiability:*

The final tally is verifiably correct

- Including:
  - Voters can check their own vote is included (*voter verifiability*)
  - Anyone can check that only authorized votes are counted, and no votes are changed during tallying (*universal verifiability*)
  - No ballot stuffing
- Sorely lacking in real-world systems

# Confidentiality

---

- No adversary can learn any more about votes than is revealed by the final tally

*Anonymity:*

The adversary cannot learn how voters vote, unless voters collude with the adversary.

- Anonymity is too weak
  - Allows vote selling and coercion of voters

# Confidentiality [2]

---

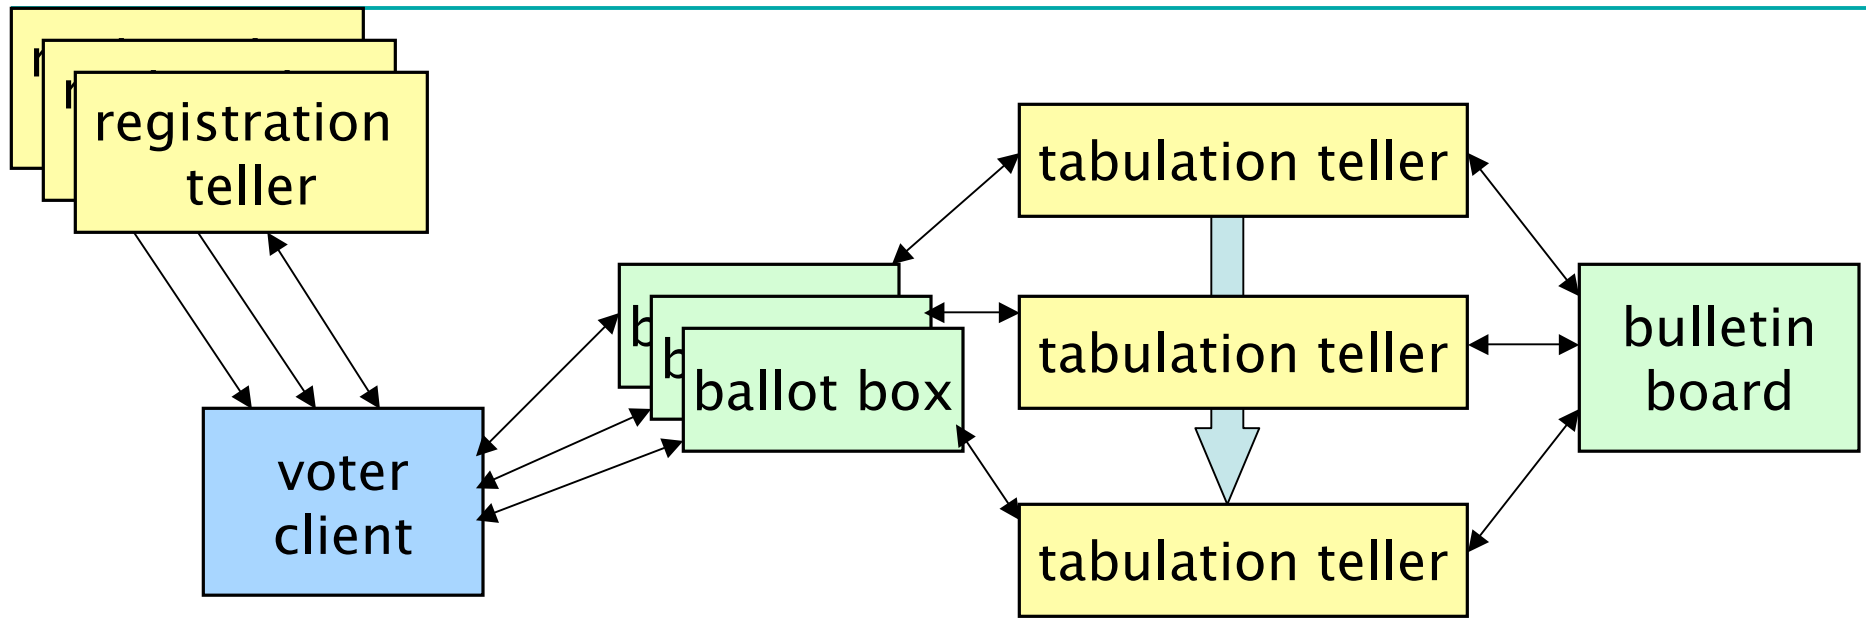
*Coercion resistance:*

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

- Required by Civitas
- Stronger than anonymity (or *receipt-freeness*)
  - Adversary could be your employer, spouse, ...
  - Must defend against forced abstention or randomization



# Civitas Architecture



## Civitas scheme

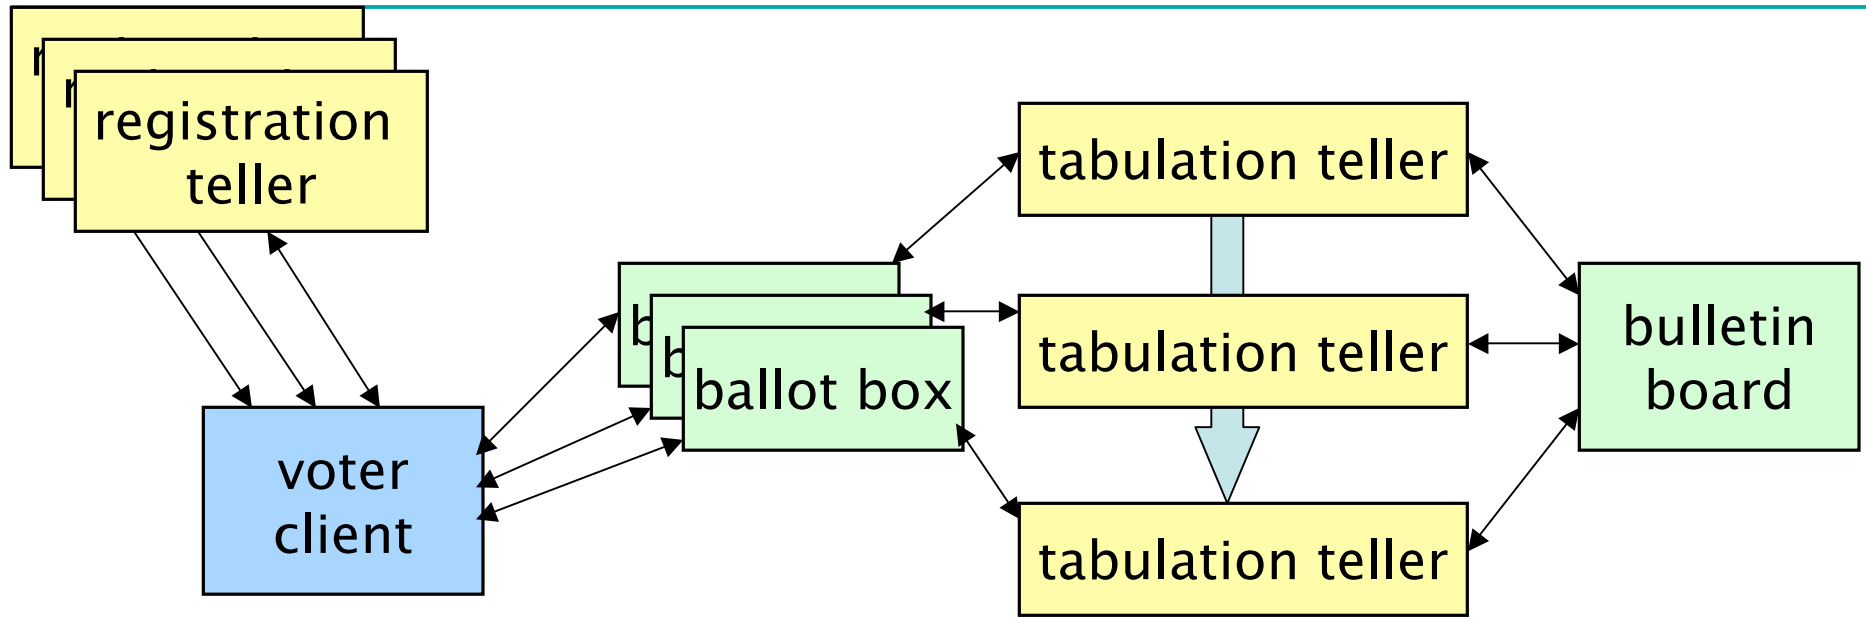
*What makes this secure? Why do we believe it is?*

# Key Idea of Scheme

---

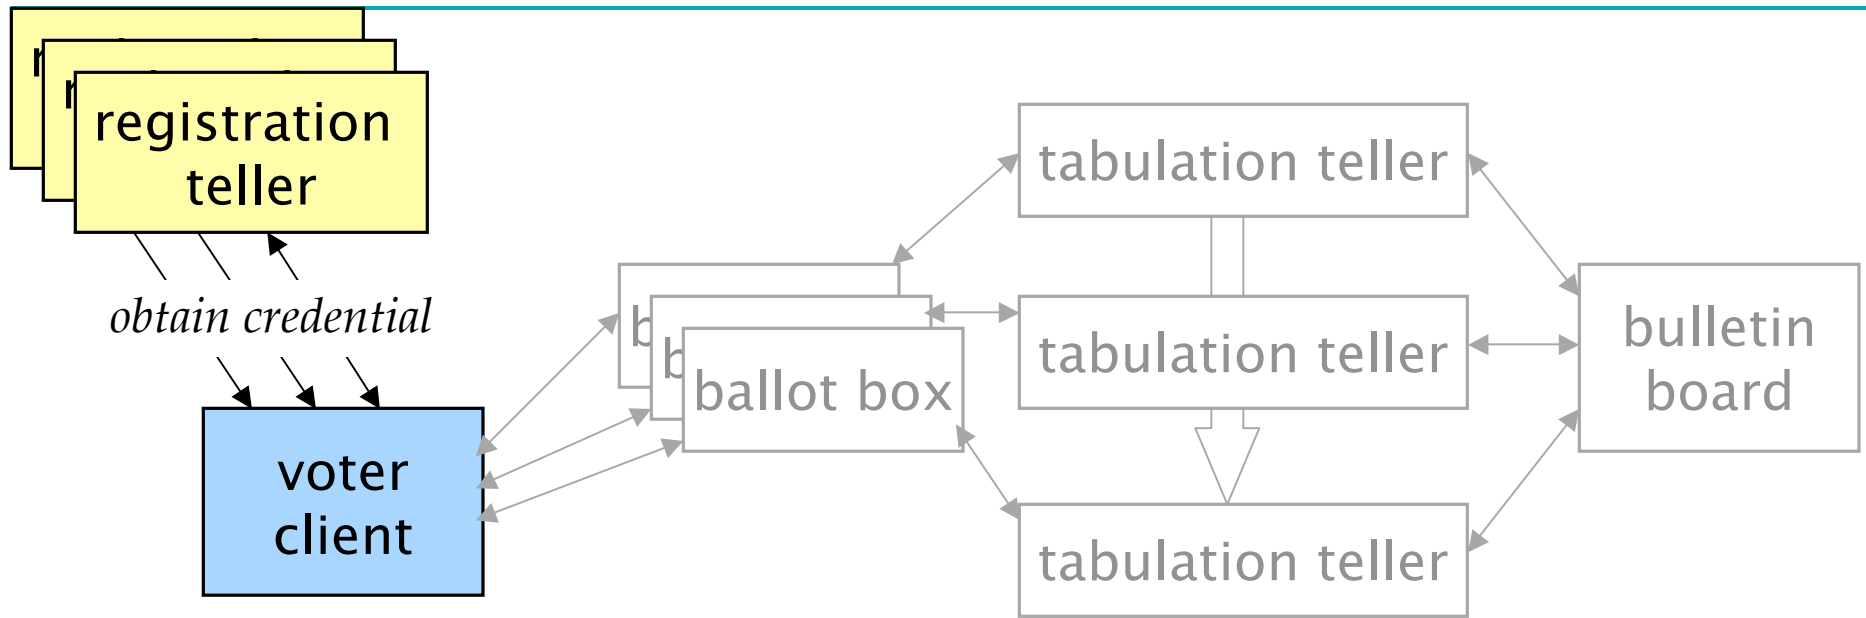
- Voter encrypts vote and “signs” with a *credential*
  - Votes are posted anonymously to ballot boxes
- Invalid credentials and votes are eliminated without revealing which were invalid
  - Tabulation tellers post ZK ("Zero Knowledge") proofs
- Anyone can verify election by checking ZK proofs
- Voter resists coercion by inventing fake credential and using it to behave exactly as coercer demands
  - Voter needs some time not under coercer’s control to use his real credential

# Cryptography



**Assumption 1.** Decision Diffie-Hellman, RSA, SHA-256 random oracle model.


# Registration



**Assumption 2.** The adversary cannot masquerade as voter during registration.

*Implement with: strong authentication,  
non-transferable secrets.*

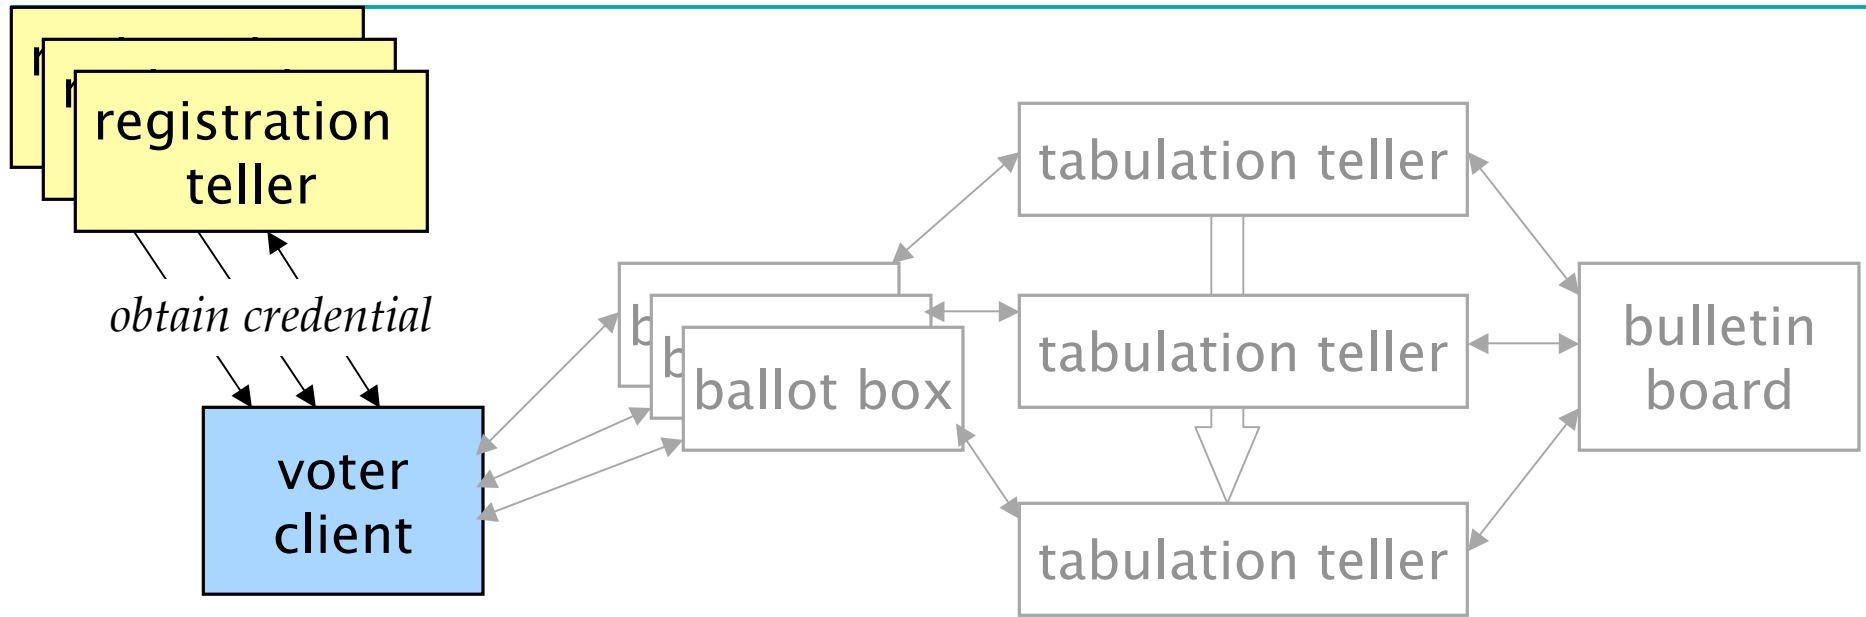
# Registration



**Assumption 3.** Each voter trusts at least one registration teller and has an untappable channel to that teller.

*Why: weakest known assumption for coercion resistance Implement with: advance, in person registration; information-theoretic encryption*

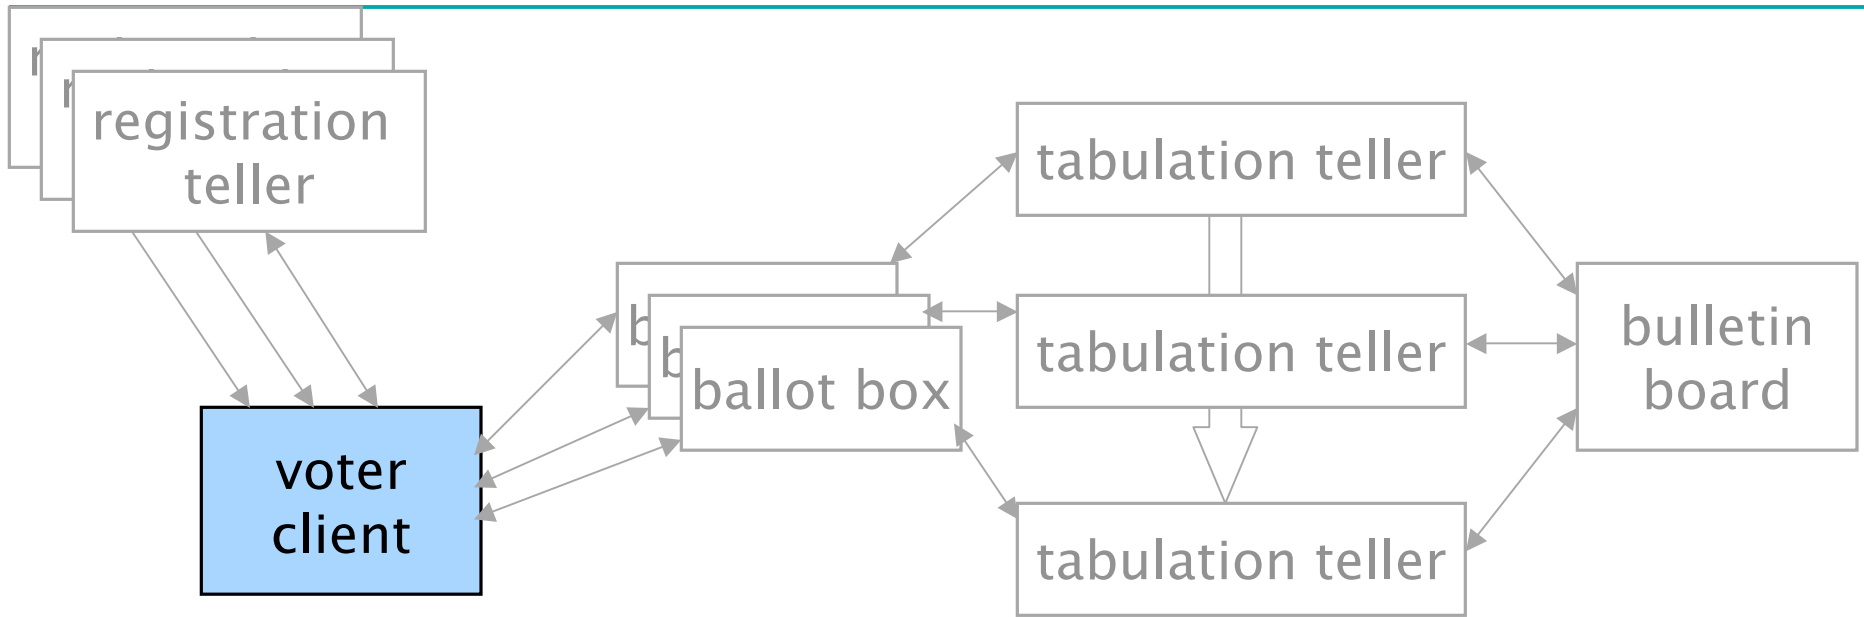
# Registration



**Assumption 3.** Each voter trusts at least one registration teller and has an untappable channel to that teller.

*Failure implies not coercion-resistant;  
doesn't impact verifiability.*


# Voting



**Assumption 4.** Voters trust their voting client.

*Reasonable: voter can choose client.*

# Voting

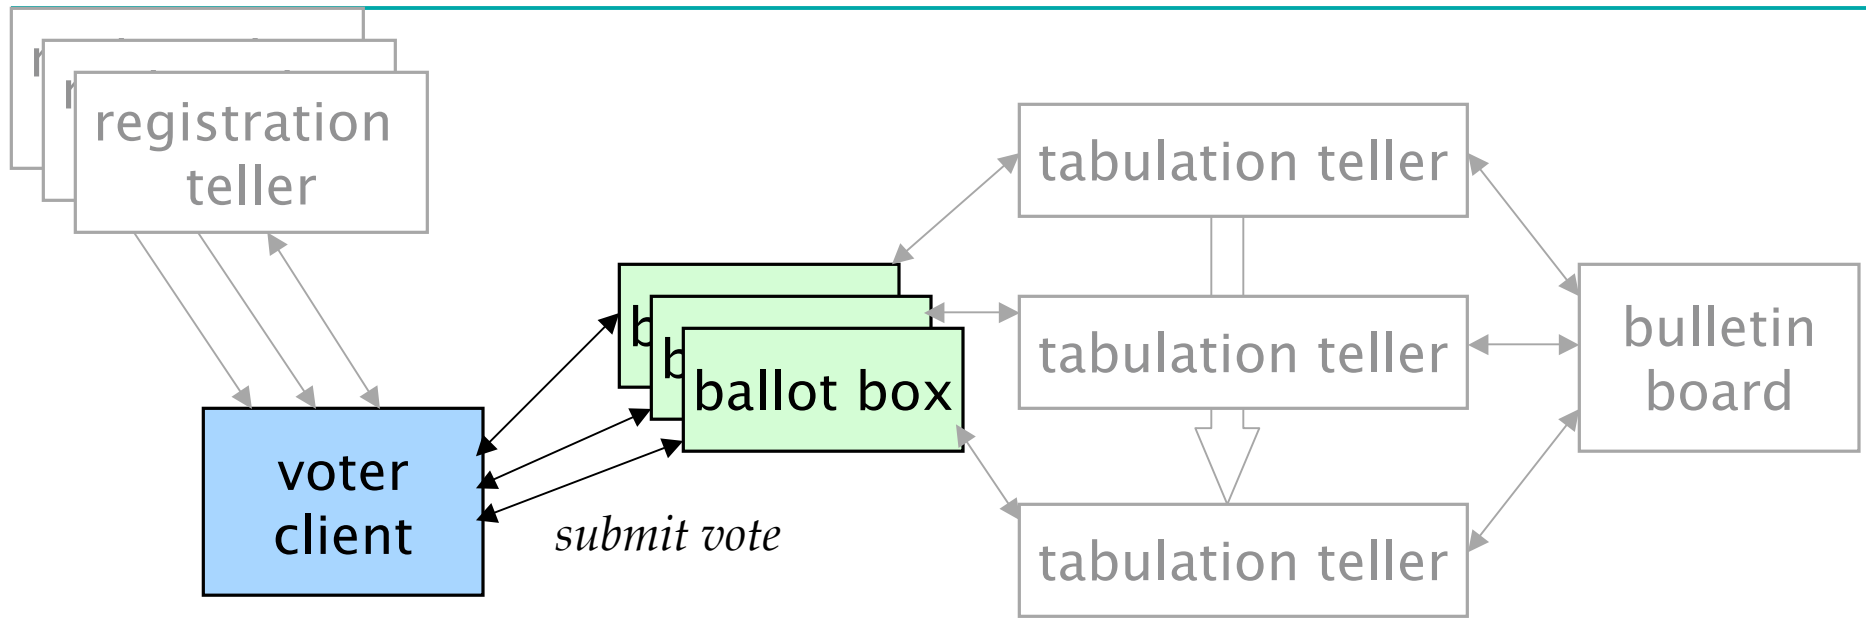


**Assumption 5.** The channels from the voter to the ballot boxes are anonymous.

*Why: otherwise coercion resistance trivially violated.  
Failure has no impact on verifiability.*



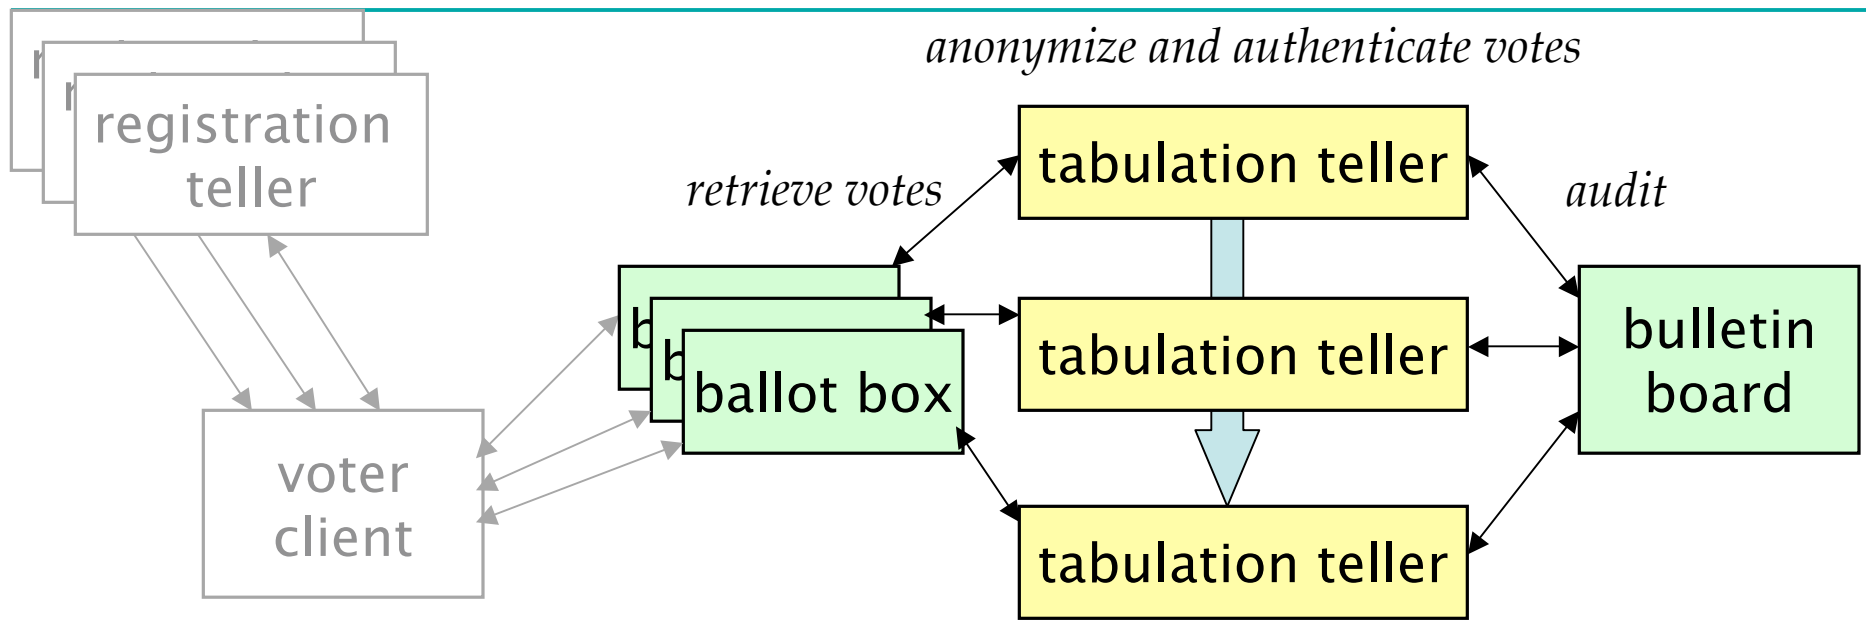
# Voting



**Assumption 6.** Each voter trusts at least one ballot box to make vote available for tallying.

*Why: expensive fault tolerance not required.*

# Tabulation



**Assumption 7.** At least one tabulation teller is honest.

*Why: keeps tellers from decrypting votes too early or cheating throughout tabulation.*

# Setup Phase

---

- Supervisor posts public keys for registrar, registration tellers, tabulation tellers
- Registrar posts identities and public keys of voters
- Tabulation tellers generate public key  $K_{TT}$  for a distributed cryptosystem
  - Everyone knows public key
  - Each teller has share of private key
  - Anyone can encrypt; participation of all tellers required to decrypt
- Registration tellers generate credentials for voters
  - Each credential is a pair of public/private values
  - Each teller responsible for generating one share of the full credential for each voter
  - Public credential share posted on bulletin board
  - Private credential share stored; later released to voter

# El Gamal Encryption

---

- $K_{TT}$  is a public key for an El Gamal cryptosystem:
  - El Gamal works similarly to RSA (same assumptions)
  - El Gamal is *malleable*: Given  $C = \text{Enc}(m, K)$  anyone can find  $D$  such that  $D \neq C$  but  $\text{Dec}(C, k) = m$
  - $\text{Enc}(m, K) * \text{Enc}(n, K) = \text{Enc}(m*n, K)$       *Homomorphic*
- The private share of a credential is  $s_i$
- The corresponding public share is  $\text{Enc}(s_i, K_{TT})$
- The complete private share is:  $s = s_1 * s_2 * \dots * s_n$
- The complete public share is  $\text{Enc}(s_1 * s_2 * \dots * s_n, K_{TT})$ 
  - The latter is computable because of the homomorphic property of El Gamal

# Voting Phase

---

- Voter retrieves private credential shares
  - Contacts each registration teller, authenticates with public key posted by registrar
  - Establishes an AES key using Needham-Schroeder-Lowe
  - Voter combines all shares to produce  $s$ , the full private credential
- Voter votes
  - Submits a copy of vote to each ballot box:

$\text{Enc}(s; K_{TT}), \text{Enc}(\text{choice}; K_{TT}), P$

- $P$  is a proof that vote is well-formed:
  - In particular, it proves that the voter had access to  $s$  and choice simultaneously

# Tabulation Phase

---

Tabulation tellers:

1. Retrieve data
2. Verify vote proofs
3. Eliminate votes with duplicate credentials
4. Anonymize votes and credentials
5. Eliminate votes with unauthorized credentials
6. Decrypt choices in remaining votes

Tellers constantly post proofs that they are performing the protocols correctly; yields verifiability

# Voters Lie to Resist Coercion

---

- Voter picks random “fake” private credential share
- Constructs new “fake” private credential
- Uses “fake” credential to behave exactly as coercer demands
  - Give it to adversary
  - Submit any vote demanded by adversary
  - Voter needs some time not under coercer’s control to use his real credential
- Yields coercion resistance

# Cryptographic Techniques

---

- Zero-knowledge (ZK) proofs
  - Vote proofs, tabulation proofs
- Plaintext equivalence test
  - Elimination of duplicate and unauthorized credentials
- Mix network
  - Anonymization



# Zero-Knowledge Proofs

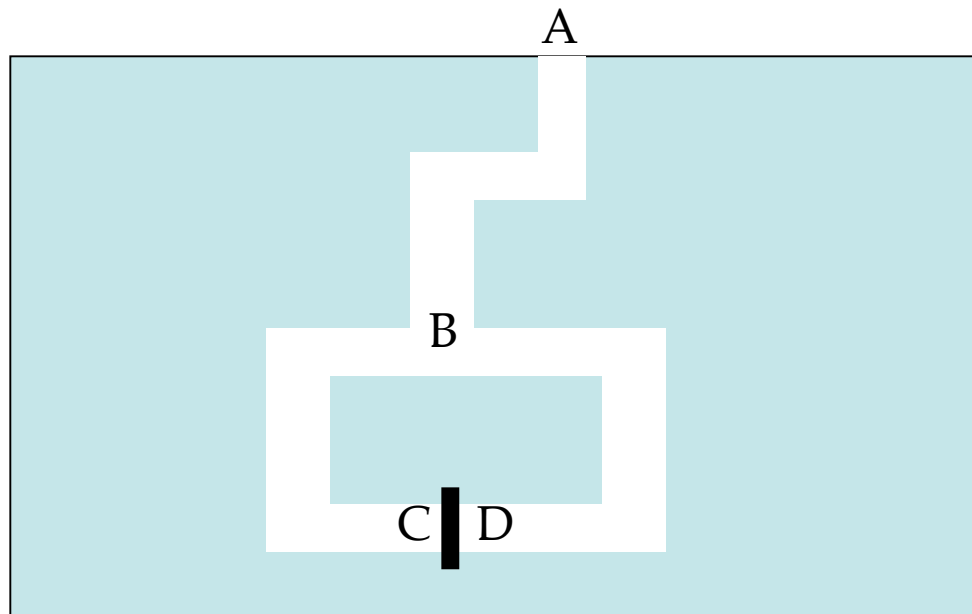
---

- Standard proofs in math class: student (prover) writes something that TA (verifier) checks.
  - Convinces verifier of statement made by prover
- But standard proofs also reveal knowledge to the verifier
  - Prover: “I know the password to the Federal Reserve”
  - Verifier: “I don’t believe you!”
  - Prover: “It’s XYZZY”.
  - Verifier: Logs into Fed to check if password works.
  - Verifier: “Thanks. Now I know it too.”

# Zero-Knowledge Proofs

---

- A *zero-knowledge* proof allows  $P$  to prove to  $V$  that a statement is true without revealing any more information.
- Example: The magic word and the cave



# Zero-Knowledge Proofs

---

- Commit:
  - V stands at A
  - P walks into cave to either C or D
  - V walks to B
- Challenge:
  - V shouts to P to either come out the left or the right passage
- Response:
  - P complies, using the magic word if necessary
- P and V repeat until V is convinced P knows the magic word


# Zero-Knowledge Proofs

---

- Zero-knowledge:
  - V never learns the magic word
  - V can't convince anyone else that P knows the magic word
    - P & V could have agreed on series of "left/right" in advance
- Large classes of problems have ZK proofs
- This proof was *interactive*
  - Based on challenge/response
  - Can make *noninteractive* by using a special kind of hash function to generate the challenge
- Plaintext Equivalence Test
  - Special kind of ZK proof
  - Tabulation tellers prove (as a group) that  $\text{Dec}(c) = \text{Dec}(c')$  without anyone, including the tellers, learning what  $\text{Dec}(c)$  or  $\text{Dec}(c')$  actually are

# Mix networks: Anonymity

- Chaum 1981: Basic Mix network
- Suppose that there are N servers with public keys  $K_1 \dots K_N$ .
- A mix message  $M_a$  looks like:  $K_1\{K_2\{\dots K_N\{m_a\}\dots\}}$
- To anonymize a set of messages  $M_1, M_2, \dots, M_j$ :
  - Server  $i$  decrypts the messages, permutes them, and forwards them to server  $i+1$
  - The last server will reveal  $m_1, m_2, \dots, m_j$  in some random permutation:



# Aside: Secret Sharing

---

- How to share a secret among  $N+1$  players:
  - Owner of the secret generates  $N$  bitstrings  $R_1 \dots R_N$
  - Player 0 gets  $S \oplus R_1 \oplus \dots \oplus R_N$
  - Player  $j > 0$  gets  $R_j$
  - All  $N$  players can cooperate to recover  $S$  -- they just XOR their shares.
- *Threshold* schemes allow  $k$ -out-of- $N$  players to recover the secret:
  - Owner of the secret picks a random polynomial  $f$  with degree  $(k-1)$  such that  $f(0) = S$
  - Player  $j > 0$  gets  $f(j)$
  - If any  $k$  players get together, they can use interpolation to calculate  $f(0)$
  - If fewer than  $k$  players get together, there's no information about  $f(0)$ .

# Blind Signatures

---

- Digital signature scheme equipped with a commutative *blinding* operation
  - Signer never learns what they signed
  - Like signing an envelope with a window (or with carbon paper)
  - I.e.:  $\text{unblind}(\text{sign}(\text{blind}(m))) = \text{sign}(m)$
- Voting scheme:
  - Voter prepares vote  $v$ , blinds, and authenticates to Authorization server, and sends vote. Server checks off voter, signs vote, and sends back to voter. Voter unblinds and now has  $\text{sign}(v)$ .
  - Voter anonymously sends  $\text{sign}(v)$  to Tabulation server. Server checks signature, then counts vote.

# Mix Networks

---

- Original Chaumian decryption mix:
  - Implemented with set of servers
  - Input: list of encrypted values
    - $\text{Enc}(\text{Enc}(\text{Enc}(\dots c \dots)))$
  - Output: same list, decrypted
    - But order of list permuted
  - Each server in mix permutes list and removes one layer of encryption
- Problem: servers trusted to be honest
  - Need *robustness*
- Civitas based on mix network
  - Actually uses *reencryption* mix



# Mix Networks

---

- Voting scheme:
  - Voter encrypts vote, authenticates to Ballot Box server, submits vote.
  - Set of tabulation tellers run a mixnet over the encrypted votes, resulting in random permutation of votes.
  - Permuted list is decrypted and tallied.

# Homomorphic Encryption

---

- A *homomorphic* encryption scheme has an operator  $\star$  such that  $\text{Enc}(m) \star \text{Enc}(n) = \text{Enc}(m \star n)$ .  $\star$  is usually either  $+$  or  $\times$ , never both.
  - E.g. both RSA and El Gamal have  $\times$ .
- Voting scheme:
  - Suppose scheme has  $+$  as homomorphism and votes are either 0 or 1.
  - Voter prepares  $\text{Enc}(0)$  or  $\text{Enc}(1)$  as vote, authenticates to Tabulation server, and submits vote.
  - Tabulation server sums all the votes, then decrypts result. Individual votes never decrypted.