


A (somewhat) gentle introduction to machine-checked cryptographic proofs

Privacy and Verifiability
for online voting

François Dupressoir

joint work with V Cortier, CC Drăgan,
B Warinschi, P-Y Strub, B Schmidt,
B Grégoire, and others



Online Voting: Some Context

- » Electronic Voting is already widely used
 - Non-governmental elections and votes
 - Legally-binding decision-making (Switzerland, Estonia, ...)

- » Problem: attacks to *privacy* and *integrity* scale up
 - Traditional approaches (observe and audit) are insufficient

- » Solution: Throw *crypto* in and mix
 - Prove strong privacy and integrity guarantees, under...
 - ... standard cryptographic assumptions, and ...
 - ... simple trust assumptions.



What is Online Voting? (Syntax)

- » The *election authority* sets up the election, ~~generates the voter roll, checks eligibility, ...~~
 - Modelled as a **Setup** algorithm
- » The *voters* cast their votes, and later may want to check them
 - Modelled as a **Vote** algorithm
- » The *ballot box* receives ballots
 - Modelled as a **Valid** algorithm (+ a chunk of state)
- » The *bulletin board* holds a public view of the ballots received and other verifiability evidence
 - Modelled as a **Publish** algorithm (+ a chunk of state)
- » The *trustees* compute the tally from ballots held within the ballot box
 - Modelled as a **Tally** algorithm
- » The *general public* may want to check the good conduct of the election
 - Modelled as a **Verify** algorithm

A Typical (secure) Online Voting System

- » **Setup** generates a keypair for the election
 - Usually shared between trustees so that a threshold of them need to collaborate to decrypt

- » **Vote** encrypts voter choice under the election public key, may protect the integrity of the ballot

- » **Valid** typically prevents direct replay of encrypted ballots, rejects ill-formed ballots
 - May prevent revotes, ...

- » **Publish** typically selects a subset of information to publish
 - May publish nothing at all (no verifiability)

- » **Verify** checks that tallying was performed correctly

Two Ways of Tallying

» Homomorphically

- **Vote** uses (partially) homomorphic encryption
- **Tally** computes homomorphically over ciphertext to get encrypted result
- A threshold of trustees decrypt the result once they agree tallying is finished, and produce a NIZK proof of correct decryption

» Using mix-nets

- **Vote** uses re-randomizable encryption
- A network of mix-servers sequentially re-randomize the ciphertexts after shuffling them, producing NIZK proofs of correct shuffling
- A threshold of trustees decrypt individual ballots once they agree shuffling is finished
- **Tally** can then be performed publicly

» We want tally-agnostic definitions for privacy and verifiability

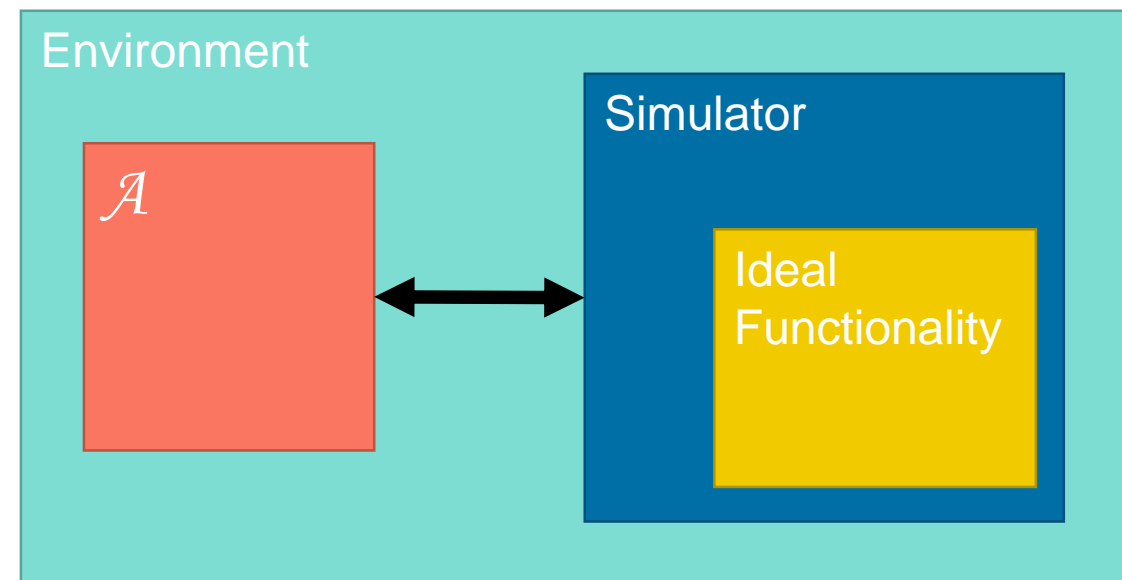
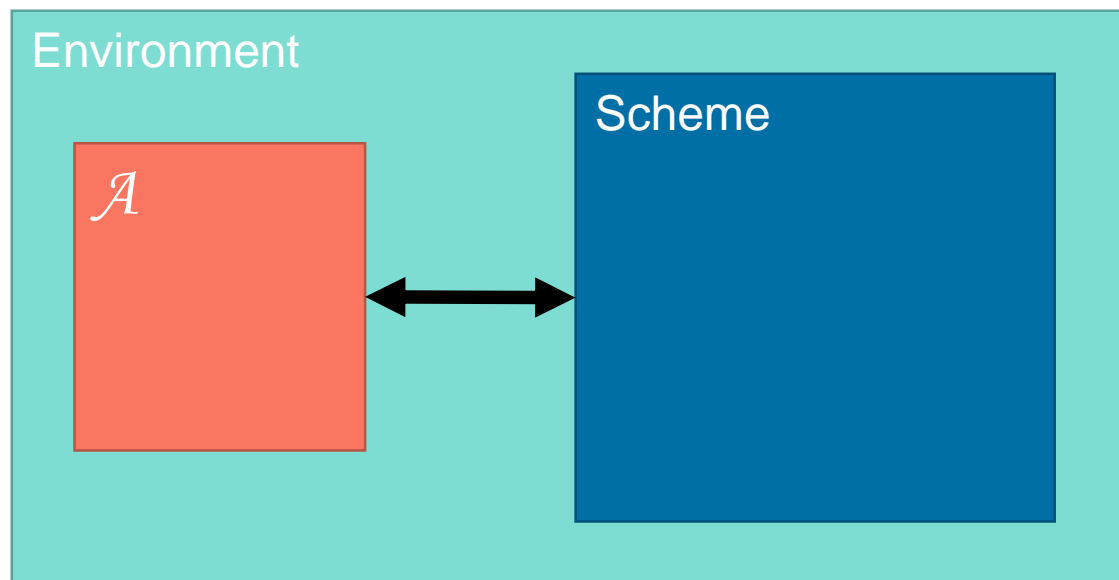


Defining and Proving Privacy for Online Voting

Take 1

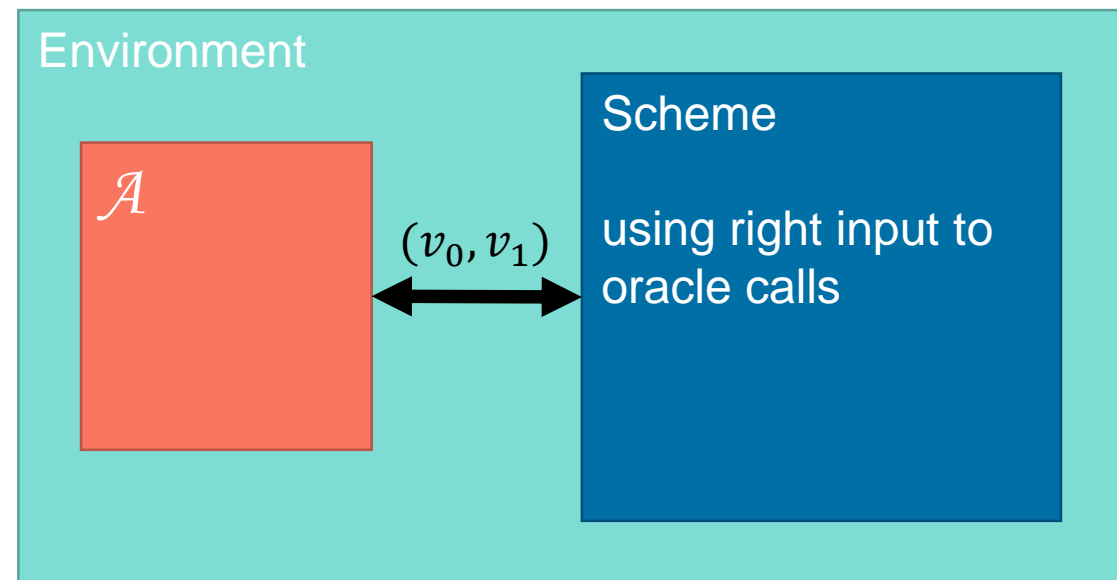
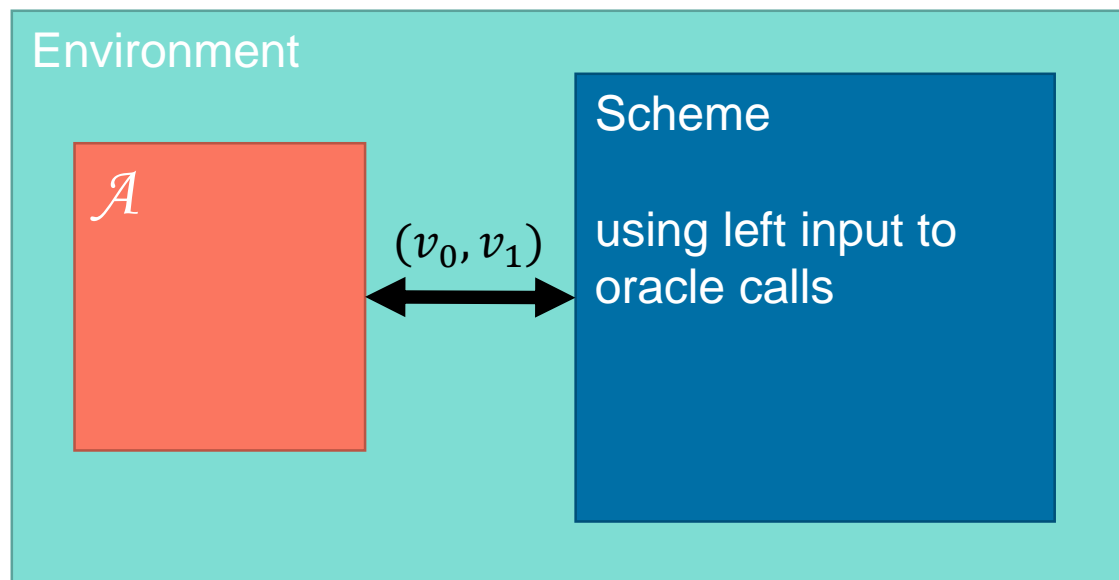
Defining Privacy I

- » Long history of bad game-based definitions
- » Ideally, want guarantees as strong as those given by a true Trusted Third Party (Simulation-Based Security)
 - There exists a simulator such that no adversary can distinguish the scheme from the simulator
- » Simulations are really hard to deal with



Defining Privacy II

- » Bernhard, Cortier, Galindo, Pereira and Warinschi (S&P 15) define BPRIV
 - Prove that, with two simpler conditions, it implies simulation-based privacy
- » BPRIV is “mostly” game-based
 - Easy to manipulate and instantiate
- » BPRIV is a “Left or Right” game



Defining Privacy III (BPRIV)

- » **Vote** oracle is the only one that is made Left or Right
 - Voter choice is the only thing whose privacy we care to protect
- » The adversary is additionally given the ability to form and cast ballots without using **Vote**
 - Models voters who may be under adversary control
- » In the Left game, everything works as expected:
 - Simply run the scheme using the left input to **Vote** oracle queries
- » In the Right game, things get ~~complicated~~ fun:
 - Maintain Ballot Boxes corresponding to both sets of inputs
 - **Publish** gives the public bulleting board produced from the Right box
 - **Tally** computes the result using the Left box,
and *simulates* a proof that the result was computed correctly from the Right box

Defining Privacy IV (BPRIV, a formal view)

$\text{Exp}_{\mathcal{A}, \mathcal{V}, \text{Sim}}^{\text{bpriv}, \beta}(\lambda, m)$

1: $\text{BB}_0, \text{BB}_1 \leftarrow []$
 2: $(\text{pk}, \text{sk}, \text{uL}) \leftarrow \text{Setup}(1^\lambda, m)$
 3: $\beta' \leftarrow \mathcal{A}^\mathcal{O}(1^\lambda, \text{pk}, \text{uL})$
 4: **return** β'

Oracle $O_{\text{cast}}(b)$

1: **if** $(\text{Valid}(\text{BB}_\beta, \text{uL}, b, \text{pk}))$ **then**
 2: $\text{BB}_0 \leftarrow \text{BB}_0 + [b]; \text{BB}_1 \leftarrow \text{BB}_1 + [b]$

Oracle $O_{\text{board}}()$

1: **return** $\text{Publish}(\text{BB}_\beta)$

Oracle $O_{\text{tally}}()$ for $\beta = 0$

1: $(r, \Pi) \leftarrow \text{Tally}(\text{BB}_0, \text{sk})$
 2:
 3: **return** (r, Π)

Oracle $O_{\text{tally}}()$ for $\beta = 1$

1: $(r, \Pi) \leftarrow \text{Tally}(\text{BB}_0, \text{sk})$
 2: $\Pi' \leftarrow \text{Sim}(\text{pk}, \text{Publish}(\text{BB}_1), r)$
 3: **return** (r, Π')

Oracle $O_{\text{vote}}(id, v_0, v_1)$

1: $\ell \leftarrow \text{uL}[id]$
 2: **if** $(\ell \neq \perp)$ **then**
 3: $b_0 \leftarrow \text{Vote}(id, v_0, \ell, \text{pk}); b_1 \leftarrow \text{Vote}(id, v_1, \ell, \text{pk})$
 4: **if** $(\text{Valid}(\text{BB}_\beta, \text{uL}, b_\beta, \text{pk}))$ **then**
 5: $\text{BB}_0 \leftarrow \text{BB}_0 + [b_0]; \text{BB}_1 \leftarrow \text{BB}_1 + [b_1]$

Proving Privacy?

- » 2 *published* attempts at proving BPRIV for Helios-like protocols
 - Both had minor issues and a significant gap

- » Zero-Knowledge proofs evidence a mathematical relation between a (secret) witness and a (public) statement

- » The language of valid statements should be in NP

- » In our case, the statement talks about the random oracle
 - If we make it stateless, not in NP
 - If we make it stateful, need new theory

- » This was never highlighted as an issue...

A New Problem

- » Nobody understands cryptographic proofs
 - Hard to write, but *even harder to read*

- » Formalize the proof in EasyCrypt
 - Introduces an asymmetry between proof writer and proof reader
 - Removes focus from the proof itself, and
 - Allows evaluator to focus on definitions and claims

- » Key insight:
 - All crypto security notions (even simulation-based) are post-conditional equivalences between open probabilistic programs
 - Relational reasoning is well suited

- » EasyCrypt allows us to dive below program logics and into semantics

You Know What we Did Last Summer (I)

» We formalize BPRIV and its associated properties for Labelled MiniVoting (Bernhard et al.)

Setup($1^\lambda, m$)

```

1: (pk, sk) ← KGen( $1^\lambda$ )
2: for i in 1..m do
3:   id ← $_s$  ID
4:   uL[id] ← Flabel(id)
5: return (pk, sk, uL)

```

Vote(id, ℓ, v, pk)

```

1: c ← Enc(pk,  $\ell, v$ )
2: return (id,  $\ell, c$ )

```

Valid(BB, uL, b, pk)

```

1: (id,  $\ell, c$ ) ← b
2:  $e_1 \leftarrow \forall id'. (id', \ell, c) \notin \text{BB}$ 
3:  $e_2 \leftarrow (\ell = \text{uL}[id])$ 
4:  $e_3 \leftarrow \text{ValidInd}(b, \text{pk})$ 
5: return ( $e_1 \wedge e_2 \wedge e_3$ )

```

Tally(BB, sk)

```

1: dbb = [ ]
2: for i in 1..|BB| do
3:   (id,  $\ell, c$ ) = BB[i]
4:   dbb[i] ← (id, Dec(sk,  $\ell, c$ ))
5: r ←  $\rho(\text{dbb})$ 
6: pbb ← Publish(BB)
7:  $\Pi \leftarrow P((\text{pk}, \text{pbb}, r), (\text{sk}, \text{BB}))$ 
8: return (r,  $\Pi$ )

```

» **Tally** is trusted, which we really do not want in practice

» We generalize over previous definitions by parameterizing the scheme

- **Flabel**, **ValidInd**, ρ , \mathcal{R}

You Know What we Did Last Summer (II)

Electric Boogaloo

» By verifiably secure refinement, we transfer the security of “Labelled MiniVoting” to:

- Mixnet-based Helios (Helios v3-mix)
- Homomorphic Helios (Helios v3-hom, Helios v4)

» By observational equivalence, we further transfer the privacy result to a previously unproved optimized version of Helios v4

» By verifiable instantiation, we obtain machine-checked privacy proofs for over 500 variants of Helios

$\text{Valid}(\text{BB}, \text{uL}, b, \text{pk})$

1: $(id, \ell, c) \leftarrow b$

2: $e_1 \leftarrow \forall id'. (id', \ell, c) \notin \text{BB}$

3: $e_2 \leftarrow (\ell = \text{uL}[id])$

4: $e_3 \leftarrow \text{ValidInd}(b, \text{pk})$

5: **return** $(e_1 \wedge e_2 \wedge e_3)$

$\text{ValidLight}(\text{BB}, \text{uL}, b, \text{pk})$

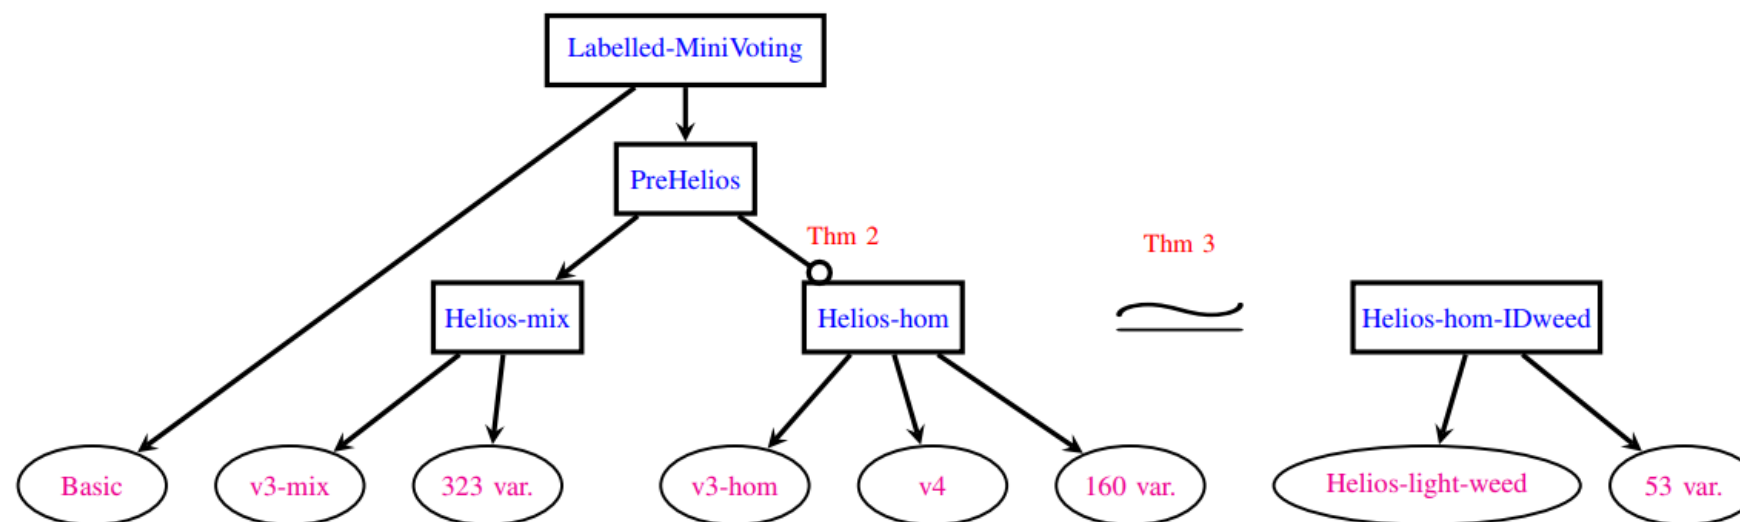
1: $(id, \ell, c) \leftarrow b$

2: $e_1 \leftarrow (b \notin \text{BB})$

3: $e_2 \leftarrow (\ell = \text{uL}[id])$

4: $e_3 \leftarrow \text{ValidInd}(b, \text{pk})$

5: **return** $(e_1 \wedge e_2 \wedge e_3)$



Verification Effort

- » About 1 person-year from start of project to “final” qed
 - Includes false starts, time for RA to learn both the crypto and formal tools
 - Roughly 1,500 lines of definitions (includes named variants)
 - Roughly 14,000 lines of proof (includes named variants)
 - Unnamed variants are automatically generated – ~150LoD, ~500LoP each

- » Initial proof for Labelled MiniVoting obtained about 75% of the effort in
 - But later iterations needed to extend treatment of proofs over relations that include random oracles

- » Identified a missing assumption in published proofs of Labelled MiniVoting
 - Does not affect practical security, since it is discharged on concrete instantiations



Defining and Proving Verifiability for Online Voting

Take 1



Verifiability in Online Voting

- » We want verifiability
 - If something goes wrong, anywhere, we want to know it
- » We want it with minimal trust assumptions
 - Force at least two parties to collude to subvert the election without detection
- » End-to-End verifiability relies on:
 - Individual Verifiability: individual voters should be able to verify that their vote was both *cast as intended*, and *recorded as cast*
 - Universal Verifiability: anyone should be able to verify that all votes were *counted as recorded*
- » But end-to-end verifiability does not prevent ballot stuffing
 - A malicious ballot box can just stick ballots in for voters who haven't voted

Belenios (Cortier et al. 2014)

- » To prevent ballot stuffing, voters need cryptographic credentials
- » A registrar manages a mapping from eligible voters to their public keys
 - The registrar does not maintain the voter roll – that *must* be trusted
- » Signed ballots are signed by voters before being cast
 - A malicious ballot box cannot stuff ballots, as it doesn't have the voters' signing keys

- » For privacy, signatures must be stripped before tallying

- » Cortier et al. provide proofs of privacy and verifiability

Machine-Checked Privacy for Belenios

» Expectation:

1. Add registration and signing in Helios definitions and proofs
2. Run EasyCrypt
3. Minimally fiddle with proof
4. Profit

» Reality:

1. Add registration and signing in Helios definitions and proofs
2. Run EasyCrypt
3. ??
4. Wat?

Privacy for Belenios: No Dynamic Corruption

» Dynamic corruption allows adversary to replay an honest ballot to learn the vote it contains

$$\beta = 0$$

$$BB_0 = [(id, b_0), (id, b_0)]$$

$$(v_0, \pi) \leftarrow \text{Tally}(BB_0)$$

Adversary sees board BB_0 ,
result v_0 , and real proof π

$$\beta = 1$$

$$BB_1 = [(id, b_1), (id, b_1)]$$

$$BB_0 = [(id, b_0), (id, b_1)]$$

$$(v_1, \pi) \leftarrow \text{Tally}(BB_0)$$

Adversary sees board BB_1 ,
result v_1 , and sim. proof π'

» Helios accidentally avoids this issue by preventing replay of ballots

- Which was put in place to stop an actual attack on privacy

Privacy for Belenios: Trust and the Registrar

- » A dishonest registrar can give invalid credentials to all voters but one
 - Tally reveals that voter's preferences

- » Is this an attack?
 - Yes

- » Should we care about it?
 - Yes

- » Why? Both attacks are artificial: the adversary truly learns nothing in practice...

Privacy for Belenios – The Big Problem

- » Current definitions and natural extensions are not robust
 - Consideration of elements usually left out of scope is needed

- » In a follow-up, Cortier and Lallemand prove that *all* current definitions of privacy imply individual verifiability

- » We can only get proofs in much weaker models than those we want

- » There may be actual attacks we are missing because of “silly” definitional issues

- » We’re *not* solving this here: we just accept a weaker definition and move on

Strong Verifiability for Belenios

- » If the adversary:
 - Controls either the registrar or the ballot box,
 - Knows the election private key, and
 - Can corrupt individual voters statically.

- » The final tally corresponds to the tally computed over:
 - The votes of all honest voters who perform individual verifiability checks,
 - A subset of the votes cast by honest voters who did not check,
 - At most as many corrupted votes as there were corrupted voters.

- » In practice, this is very strong as it gets
 - Gives precise bounds on the distance between final result and actual result
 - The adversary does not know who will check; statistical arguments can give tighter bounds

Machine-Checked Verifiability for Belenios

- » Really nothing to report; some extensions and clarifications to Cortier et al.'s result:
 - Give the private election key to the adversary
 - Refine what it means for a voter to have checked their ballot (in case of revotes)
 - When registrar is dishonest, even honestly generated ballots may be invalid and cannot be counted

- » The proofs are “straightforward formalizations”

- » Could further refine to allow checking on intermediate bulleting boards hen they are published

Belenios	LoC	Ver. Time (s)	Code Sim. (%)	Dev. Effort (PW)
General Concepts	5936	348	55% Helios	4
Privacy	2700	238	75% Helios	2
Verifiability	14590	1523	-	20
Variants	47030	3965	95% Belenios	1



Next Steps in Machine-Checked Cryptographic Proofs



Take 1

Go Down from the Bottom

- » Our proofs do not cover the primitives:
 - Mix-nets are assumed to be perfect obliviously permuting decryption oracles
 - Zero-Knowledge proofs are taken as assumptions
 - Including those whose statements talk about random oracles
 - Encrypt+PoK is taken as non-malleable encryption

- » These proofs are fun
 - Fun number theory
 - Interesting proof techniques, where simulators can rerun adversaries with fixed randomness
 - Zero-Knowledge is still not very well understood in terms of proofs, composition
 - Even though these things are conceptually simple, they involve *interactive systems*

Go Up from the Top

- » Interactive systems are also increasingly used by the crypto community for compositional security
 - Constructive Cryptography
 - Universal Composability

- » Having proof tools that support them will be crucial in scaling machine-checked crypto up to larger constructions

- » Ideas from distributed system verification could be looked into

Summary

- » Machine-checked crypto is costly
- » But worth it for select applications where trust in the system is paramount
 - Standards, voting, e-government, privacy, ...

- » Definitions of privacy for electronic voting are brittle and inadequate

- » PL and PV can still contribute to machine-checked crypto
 - Weird rewinding semantics
 - (Relational) Semantics for interactive open probabilistic programs



UNIVERSITY OF
SURREY