Risk Limiting Tallies and Verification

Wojciech Jamroga, Peter B Roenne, Peter Y A Ryan, Philip B Stark

Université du Luxembourg &
Centre for Security and Trust (SnT)
University of California, Berkeley

Evote-Id 2019

Outline

- Motivation
- Sketch of E2E verifiability
- Sketch of risk-limiting audits
- Risk-limiting tallies
- Sketch of Selene
- Risk-limiting verifiability

Motivation

- Even a scheme satisfying all the usual privacy, coercion-resistance properties may fail to provide vote privacy in some corner cases, e.g. unanimous vote, no votes for X etc.
- Suffices that this is perceived as a possibility.
- Also threat of “Italian” aka signature attacks.
- And the “sting-in-the-tail” in Selene.

Commune	Polling station	Grzegorz RA	Grzegorz Jar	Karol Roma	Aleksandra I	Jan KILIAN	Tadeusz CY	Anna Magda	Ewelina Mar	Dorota DUD	Marian Czes
gm. Przywidz	Remiza Strażacka	7	7	6	0	0	10	1	1	6	2
gm. Przywidz	Gimnazjum	23	3	11	4	0	23	4	0	6	0
gm. Przywidz	Szkoła Podstawowa (Trze	10	5	5	3	0	7	4	0	6	1
gm. Pszczółki	Urząd Gminy w Pszczółka	4	3	6	5	4	16	3	1	1	1
gm. Pszczółki	Szkoła Podstawowa w Psz	14	1	2	1	0	20	1	2	2	0
gm. Pszczółki	Szkoła Podstawowa w Sk	2	2	1	4	3	20	0	1	1	0
gm. Pszczółki	Szkoła Podstawowa w Żel	2	0	2	2	0	5	1	1	2	2
gm. Pszczółki	Szkoła Podstawowa w Ró	9	0	4	6	1	21	1	1	2	3
gm. Pszczółki	Publicznme Gimnazjum w	4	1	3	3	0	22	0	0	3	1
gm. Pszczółki	Fundacja „Żyć godnie" K	3	0	0	1	1	7	0	0	4	0
gm. Suchy Dąb	Zespół Szkół	5	3	4	4	1	9	3	0	2	0
gm. Suchy Dąb	świetlica wiejska	4	0	3	6	0	9	0	0	1	0
gm. Suchy Dąb	Zespół Szkolno-Przedszk	1	0	3	2	1	6	1	1	1	0
gm. Trąbki Wielkie	Szkoła Podstawowa w Cz	4	3	9	4	5	12	1	1	3	5
gm. Trąbki Wielkie	Szkoła Podstawowa w Mi	8	2	13	9	0	12	2	1	2	0
gm. Trąbki Wielkie	Szkoła Podstawowa w Sol	5	3	46	10	3	10	2	2	2	0
gm. Trąbki Wielkie	Szkoła Podstawowa w Trą	6	0	71	4	3	22	2	2	5	0
gm. Trąbki Wielkie	Szkoła Podstawowa w Kł	4	0	19	1	2	14	0	2	1	0
m. Kwidzyn	Budynek Zakładu Utyliza	18	2	2	2	3	41	5	0	3	0
m. Kwidzyn	Przedszkole Niepubliczne	14	2	2	4	3	26	1	0	5	2
m. Kwidzyn	Przedszkole Niepubliczne	19	6	2	6	4	34	3	1	7	1
m. Kwidzyn	Zespół Szkół Ogólnokszta	16	1	1	6	3	14	4	0	3	1
m. Kwidzyn	Centrum Kształcenia Zaw	15	1	2	4	12	28	0	1	2	0
m. Kwidzyn	Szkoła Podstawowa Nr 2	28	3	1	2	7	37	6	0	3	0

Motivation II

- Typically just accepted as a fact of life, but maybe we can do better.
- Tally hiding schemes help, but are computationally intensive and arguably lack transparency.

Key Idea

- To apply risk-limiting techniques, but now applied to the tally rather than the audit.
- Reveal sufficient votes, randomly selected, to achieve the required confidence level, e.g. 95%, leaving a proportion unrevealed.
- Provides plausible deniability: voter just claims that the required vote must be amongst those shrouded.
- Can be applied to any E2E V scheme involving posting the encrypted votes to a BB, e.g. Pret a Voter, Helios, PGD, Selene, etc.

E2E Voter-Verifiability

- Goal: voters can confirm that their vote is accurately counted (while avoiding coercion, vote-buying etc).
  - At the time of casting voters get a "receipt"; an encrypted/ encoded representation of their vote.
  - Cast, encrypted votes are posted to a secure, public bulletin board (ledger). Voters can verify that their receipt is correctly posted.
  - A (universally) verifiable, anonymising tabulation is performed on the posted receipts.

Public Bulletin Board

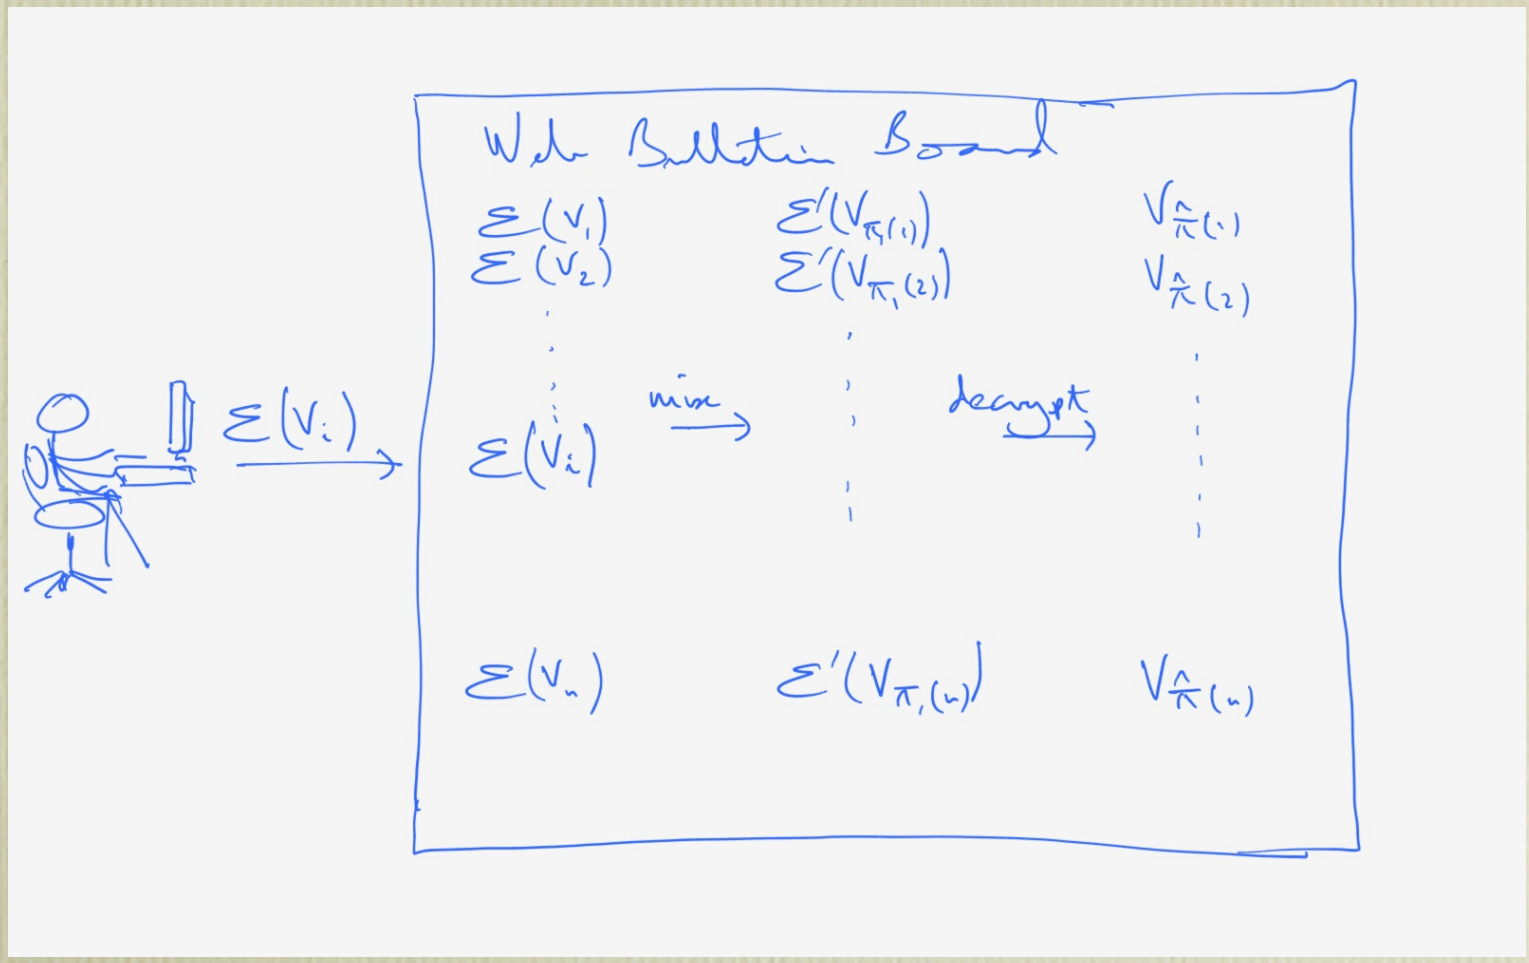

Risk-Limiting Audits

- Due to Philip Stark (UCB).
- Typically used to provide assurance in a e-tally.
- Assume a well-curated paper audit trail.
- Random sampling to develop confidence in the hypothesis: the outcome, i.e. the winners(s).
- Continue sampling until the required confidence level is achieved or a full hand tally (which replaces the original outcome).

Risk-Limiting Audits II

- The maximal chance that a wrong outcome will be accepted is the *risk limit*.
- *Comparison* audits where a link exists between the paper and digital tally of each individual ballot or batch of ballots.
- Otherwise *ballot-polling*.

Risk Limiting Tallies

- We just need a good E2E V scheme that posts to the (shuffled) encrypted ballots to the BB.
- We will perform ballot polling RL: select a random subset of the $\{V_i\}$, decrypt these and compute the risk-limit and extend the sample as necessary.
- Think: sampling from L to R from a random permutation.
- We can also sample with replacement by reshuffling between samples.

Sample Sizes Near Unanimity

candidates	α								
	10^{-1}	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}	10^{-9}
2	5	9	13	17	21	24	28	31	35
10	9	13	17	20	24	27	31	34	38

Certifiable Random Sources

- We need “good” sources of randomness, not just unpredictable but also “God-given”.
- Various possibilities:
  - According to the output of the mixes.
  - Public ceremony with dice or lottery.
  - Beacon, e.g. NIST.
  - Stockmarket values.
  - Algorand style: commitments plus Verified Random Functions… etc…..

RL without null hypothesis

- A complication is that we don't have a null hypothesis.
- But Philip has solved this one (see paper), and where the hypothesis is just the winner(s), but need larger samples.
- Results of independent interest.
- We may be able to supply a null-hypothesis (the winner(s)) based on a secret tally by trustees.

Narrow margins etc

- If winning margins are narrow, RL techniques may result in (almost) all ballots being revealed, undermining the plausible deniability goal.
- At first glance it seems that narrow margins should not be a problem, but in some cases it might: e.g. A and B in close tie and X very unpopular.
- A number of strategies are available to handle this:

Plausible deniability strategies

- If it starts looking like a close race between A and B we can start PETs of further {v_i}s against {A} and {B}.
- Or we switch to tally-hiding, essentially MPC.
- We could decide the strategy based on a secret tally (need to be careful what we leak here).
- In any case we can guarantee say >=10% of ballots stay shrouded.

Part 2:
Risk-Limiting Verification

Selene

- A very simple approach to E2E V: give each voter a private tracker number and post these on the WBB alongside the vote in the clear.
- Verification is simple and intuitive-no need for voters to handle encrypted ballots etc.
- But obvious problems, including tracker collisions and coercion.

Tracker numbers

347563	Obelix
947253	Asterix
556884	Panoramix
569331	Idefix
586994	Idefix
607855	Obelix
374823	Obelix

The goals of Selene

- To guarantee that each voter is assigned a unique tracker number.
- To notify the voters of their trackers (after trackers/votes pairs have been posted) in a way that provides **high assurance that it is "correct", i.e. unique, but is deniable.**
- And we do this in a way that ensures no single entity knows the assignment.

The Setup

- For each voter we post to the WBB:
- $PK_i, \{n_i\}_{PK\_T}, TDC_i\{n_i\}$
- $\{n_i\}_{PK}$ will be used in the tallying.
- $TDC_i\{n_i\}$, Trap Door Commitment for voter i, will be used in notifying the voter of the tracker.
- $PK_i, \{n_i\}_{PK}, [g^{r\_i}], g^{n\_i} \cdot h_i^{r\_i}$

Notifying the trackers

- Trustees reveal g^{r_i} to the i-th voter through a private (untappable) channel.
- The voter can now pair this with the TDC to form the ElGamal cryptogram:
- $$(g^{r\_i},\ g^{n\_i} \cdot h_i^{r\_i})$$
- which she can decrypt as usual with her secret key x_i to reveal: n_i.

Coercion Mitigation

- If V_i is coerced she can compute, with knowledge of the trapdoor, an alternative $(g^{r_i})'$ value which will open the encryption to a tracker number to satisfy the coercer.
- On the other hand, without the knowledge of secret trapdoor, this is intractable, so an attacker cannot reveal the wrong tracker to the voter.
- Sort of magic bank deposit box.

The sting in the tail!

- A coerced voter might by mischance chose the coercer's tracker.
- Or, the coercer simply claims that it is his tracker number anyway.
- Or he coerces many voters and we get collisions.
- Some variants of Selene to address this, but typically loose transparency.

Risk-Limiting Verification

- RL techniques can help here too: not reveal all the trackers.
- Reveal just the trackers associated with revealed ballots?
- Note: can run RLV independent of any RLT.
- But do we notify voters of unrevealed trackers? Seems dangerous not to.

Nice, but....

- But the coercer could still demand the voter to reveal his tracker, and then again claim that it is his.
- To mitigate this we could avoid revealing the set of assigned (valid) trackers, but voters need to know if the revealed tracker is valid.
- Could just draw them from subset with negligible cardinality, e.g. six digits, or publish an excess number etc.
- Coercion resistance authority?

Discussion

- Are we side-stepping a (hitherto undiscovered?) impossibility result by relaxing the properties and introducing a probabilistic component?
- BTW, reminiscent of Ron's distinguishing example for coercion vs vote-buying: voter gets a (plaintext) receipt with 50% probability.
- Compare also Random Sample Voting.

Conclusions

- Risk-limiting techniques applied to the tallying improves coercion resistance, while retaining appropriate confidence levels.
- But is it “undemocratic”?
- Also improved coercion mitigation when applied to the verification steps, in particular for Selene.
- Not so clear for general E2E V schemes: presumably need a verifiable, random allocation of ballot receipts to the voters.

Thank you!