

# An Introduction to Punchscan

Stefan Popoveniuc, Ben Hosp  
George Washington University - CS Dept.  
Washington DC 20052  
{poste,bhosp}@gwu.edu

October 15, 2006

## Abstract

Punchscan is a precinct-read optical-scan balloting system that allows voters to take their ballot with them after scanning. This does not violate the secret ballot principle because the ballots cannot be read without secret information held by the authority in charge of the election. In fact, this election authority will publish the ballots for everyone to see, allowing voters whose ballots were incorrectly omitted to complain. Punchscan vote-counting is performed in private by the Election Authority – who uses their secret information to decode the ballots – but is verified in public by an auditor.

In this paper we describe how and why PunchScan works. We have kept most of the description at an outline level so that it may be used as a straw model of a cryptographic voting system, and so that our paper satisfies the page limits of VSRW06. Section 6 presents a summary of the technical details; for a full technical paper about PunchScan, please refer to <http://home.gwu.edu/~bhosp/punchscan/article.pdf> As this full technical paper is not in review at any other venue, if the VSRW06 Technical Committee wishes, this latter document may also be used at VSRW06.

## 1 Motivation

The accurate results of a democratic election are at the heart of any modern society. Democracies are built throughout the world with the commitment to have elected individuals representing the entire population of a nation. To be able to record the wish of the people accurately we need to have a voting system that is transparent, reliable and verifiable. We need to be able to prove that the elections are run correctly, that every vote counts, and that the every person going to the polls and exercising their right to vote can make a difference. At the same time, we have to respect the secret nature of any vote. Linking a voter to a vote should not be possible, with or without the complicity of the voter.

PunchScan is a novel voting system and extremely easy to use, both by the voter and by the people running the elections. It is transparent and reliable, and provides public verifiability, election integrity and enhanced voter privacy.

## 2 Key elements/Ideas

There are two key elements that make PunchScan work:

1. The ballot is made out of two separate pages. When the two pages are put together, the resulting ballot reveals the choices of the voter. When only one page is viewed, it gives no information – in the computational sense – about what candidates the voter chose. Thus, if one page of the ballot is destroyed, the voter can keep the other page, without violating ballot secrecy.
2. A mechanism that allows the recovery of the candidate choices from only one page of the ballot
3. A proof of the integrity of the election, through pre and post election audits.

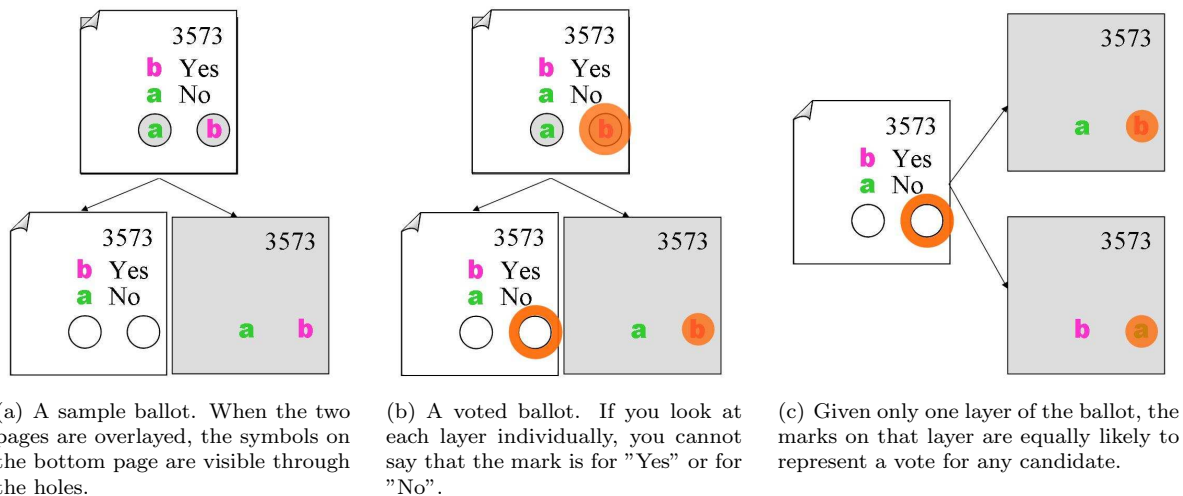


Figure 1: PunchScan's ballot

These ideas are common both to PunchScan, and to a previous method of David Chaum's [?]. However, PunchScan is more practical, because it does not suffer from the perfect alignment problem of the previous method, because the cryptography used is simpler, and because the time required to find the result and obtain the integrity proof is smaller.

### 3 High-level system design

PunchScan achieves publicly verifiable integrity by a simple blend of paper voting and optical scan. It gives each voter the opportunity to take their vote home and check that it is counted in the final tally. In this section, we first describe the ballot itself, then we present all the phases of the voting process as seen by all the participants: voters, the election authority, and candidates.

We assume that the candidates are auditing the election, since they are the ones that should care most about a correct outcome; in particular, each candidate would want to check that his rightful votes were not given to another candidate.

#### 3.1 Ballot design

A ballot consists of a single sheet of paper, transversally folded in half to form two pages. The top page of the ballot has holes in it, and the information on the bottom page can be read through the holes. The top page also contains all the text needed on the ballot, such as contests (i.e.: ballot questions) and the candidates' names. Every answer has a symbol assigned to it and the assignment of symbols to answers varies from ballot to ballot. On the bottom page of the ballot, there is an (apparently) unordered list of symbols and their order differs from ballot to ballot. The top and the bottom ballot pages are aligned in such a way that when they are overlaid, for every question on the ballot, the symbols from the bottom page are visible through the holes made on the top page (see figure 1(a)).

In PunchScan, the voter uses a dauber to mark the selection of candidates. A dauber is a pen that leaves a disk of ink on the paper when it makes contact, just like the ones used by Bingo players to mark the numbers on their tickets. The diameter of the ink disc is greater than the diameter of the hole punched through the top page, which means the dauber leaves a mark on both the top and bottom ballot pages. Figure 1(b) contains a ballot voted for "Yes".

Because the order of the symbols on the two pages of a ballot is different (and independent), one cannot determine which mark is for which candidate by viewing only one page. We assume that the association of candidates with symbols and the order of the symbols on the bottom page are uniformly random. Figure 1(c) has the right answer selected on the top layer; depending on which possible bottom layer is this ballot's

actual bottom layer, that mark could represent a vote for "Yes" or a vote for "No", both with a probability of 50%.

## 3.2 Chronological description

There are three phases of the voting:

- the preelection phase (labeled B for *Before*)
- the election day (labeled E for *Election*)
- the post election phase (labeled A for *After*)

### 3.2.1 The preelection phase

The preelection phase is a preparatory one, allowing the setup of the election and allowing integrity proofs to be carried out. During the preelection phase, the ballots are generated, printed and audited. Also, the information that allows recovering the choice from one page of the ballot is generated and checked. The chronological order is the following:

- B.1 The election authority generates ballots and commits to them.
- B.2 The election authority generates and commits to the information necessary for decrypting one page of the ballot when the other one is destroyed.
- B.3 The candidates challenge the election authority and ask to see some of the ballots (say half), along with the information from B.2.
- B.4 The election authority provides the requested ballots, and opens the commitments associated with them, thus spoiling them.
- B.5 The candidates check to ensure that the commitments are consistent with the opened ballots.

### 3.2.2 Election day

On election day, the voters go to their assigned polling places, authenticate themselves as legitimate voters, and get a ballot from the election officials.

- E.1 The voter choose either the top or bottom page of her two-page ballot, and informs pollworkers of her choice.
- E.2 In a private voting booth, the voter marks her favorite candidates on the ballot.
- E.3 The voter separates the two pages of the ballot. The chosen page from E.1 is scanned and recorded as a cast ballot (all future references to "ballot" will refer to this surviving page), and the other page is destroyed.
- E.4 The voter can now take the scanned page home; it is her receipt. It cannot be used by the voter to convince anyone of how she voted, but she can check to make sure that her ballot is eventually counted.

Note that because the voter must commit to a choice of ballot layer before voting, several problems common to paper-ballot based voting systems (such as chain voting) are prevented.

### 3.2.3 The post election phase

After all the polls close, the election is audited and proofs carried out to ensure the integrity of the election. The chronological order of the events following an election is as follows:

- A.1 Any voter can go to the election authority web site, enter a serial number for her ballot, check that the ballot is there and that it accurately resembles the page she possesses.
- A.2 The election authority processes all ballots to produce decrypted versions, along with an intermediary (partially decrypted) form of all the ballots
- A.3 The candidates ask to see some of the transformations from the original ballots to the intermediary forms, and some of the transformation from the intermediary form to the clear form.
- A.4 The election authority replies to the challenges made by the candidates in A.3
- A.5 The candidates check to see if the reply of the election authority is consistent with the commitments made in the preelection phase [B.2], and with the information made public in [A.2].

## 4 Description by roles

### 4.1 The voter

On Election Day, the voter comes to the assigned polling place and authenticates herself as a legitimate voter. She gets a ballot and a dauber and enters a private polling booth. She chooses her favorite candidates by making a mark with the dauber on the symbol associated with the candidate. She then shreds one of the pages of the ballot, and keeps the other one. Then, she scans this page. She may walk out of the polling place with this page, which serves as her (encrypted) receipt. Later, she can go to a web site, type in the serial number of her ballot and check that the ballot is there. No other checks are required from the voter.

### 4.2 The election authority (EA)

In the preelection phase, EA decides the format of a canonical ballot. This is the one from which all the other ballot variants will be generated. Also, the canonical ballot is used to recover the choices of the voters, after one page of the ballot has been destroyed.

EA generates at least twice the number of ballots needed in the election, and commits to them (making the commitment public; the ballots themselves remain secret). It also generates and commits to information necessary to recover the intent of a voter from one page of the ballot.

In response to the preelection challenge, [B.3], EA discloses all the information about half of the ballots (thus spoiling them). This allows the candidates to check the commitments and ensures, with high probability, that all the ballots have been correctly generated.

After the election, EA posts partially decrypted ballots and clear text ballots. To prove that the decryption (the partial one and the final one) was done correctly, for each vote EA will reveal either how it transformed the voted ballot into a partially decrypted one, or how it transformed a partial decrypted ballot into a clear text one, but not both for the same ballot. The auditors choose which part will be revealed, and the chances of a cheating EA being detected grow exponentially with the number of votes cheated on.

### 4.3 The candidates

We assume that the candidates are competing in an election. Because of this, we can safely allow the candidates also to play the role of auditors. As auditors, the candidates challenge EA during preelection and post election and check that the replies are consistent with the commitments.

$\oplus$	Non Flipped	Flipped
Non Flipped	Non Flipped	Flipped
Flipped	Flipped	Non Flipped

Table 1: Flipped / Non Flipped logic

## 5 An Example

We describe a minimal example: the election consists of a single binary contest; the voters vote “Yes” or “No”. The EA decides that, in the canonical ballot, the symbol “a” is associated with “Yes” and the symbol “b” with “No” on the top page. The EA also decides that the order is “a” “b” on the bottom page. The canonical ballot is presented in figure 2(a). The EA defines what is a shift of one from the canonical form on top and bottom pages. The canonical ballot corresponds to a shift of 0 (call it a non-flipped ballot) and the non-canonical ballot corresponds to a shift of one (call it a flipped ballot). Figure 3(a) contains all the possible top and bottom pages. Any top page can be combined with any bottom page to give out a ballot as seen in figure 3(b). The four types of ballots are equally likely.




Figure 2: The canonical ballot on a Yes/No contest

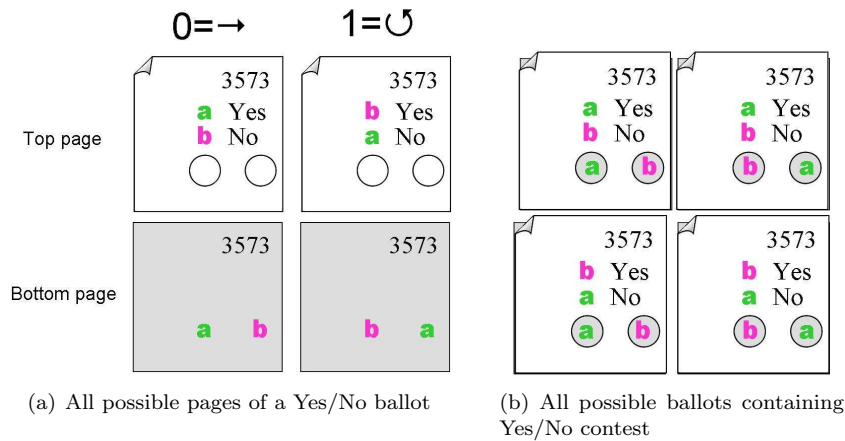


Figure 3: PunchScan’s ballot

A non-flipped top page combined with a flipped bottom page results in a flipped ballot. All the possibilities are in table 1. Note that we are only interested in knowing if the entire ballot is flipped or not, not individual pages.

To decrypt one page of the ballot, it is necessary to know if it came from a flipped or non-flipped ballot, to know if it should be flipped or not to get the canonical ballot. In Punchscan, this information is split into two flip/non-flip operations (flip1 and flip2) for each ballot. These operations, combined, will take the ballot page to the canonical ballot. The information is split so that one half of the split information can be

made public for auditing purposes. The relation that has to hold between the pages of the ballot and the information used for recovering is:  $\text{top} + \text{bottom} = \text{flip1} + \text{flip2}$ .

The EA makes public commitments to the ballots and to flip1 and flip2. The candidates choose half the ballots at random and the election authority makes public the requested ballots along with the flip1 and flip2 for each ballot. Anyone can check that the equation  $\text{top} + \text{bottom} = \text{flip1} + \text{flip2}$  holds. Only the ballots that were not made public in this phase (pre election) will be further used in the election.

During the election phase, the EA publishes all the marked pages (half ballots) as voted on by voters. After the election, it publishes the intermediary state of the ballots (ballots + flip1) and the decrypted ballots (ballots + flip1 + flip2). These are commitments to the values of flip1 and flip2 used in the decryption of the voted half ballots.

Post election, the EA is asked to open either flip1 or flip2, but not both, since opening both would allow the linking of a voted ballot to the corresponding decrypted one. Also, it is necessary that the intermediary state ballots and the decrypted ones be in a random order (distinct from the order of the voted ballots).

The election authority defines the following tables:

- $P$  (for **Print**)
- $D$  (for **Decrypt**)
- $R$  (for **Results**)

The  $P$  table is indexed by ballot serial number and contains the top page ( $P_1$ ), bottom page ( $P_2$ ), and space for the filled-in vote (to be entered after the election). It also contains commitments to  $P_1$  and  $P_2$ .

The  $D$  table contains the first ( $D_2$ ) and second ( $D_4$ ) mark permutations (flips), the intermediary vote ( $D_3$ ) to be filled in during decryption, and information to connect it with the  $P$  table ( $D_1$ ) and the  $R$  table ( $D_5$ ). It also contains a commitment for each row of  $D$ , as well as a commitment for columns  $D_1$  and  $D_2$ , and another commitment for columns  $D_4$  and  $D_5$ .

The  $R$  table contains the clear text votes (after post-election decryption).

Consider further, for the purposes of illustration, an election with only six votes. The clear data in all the tables is in table 2. Before the election, but after the EA made the commitments, the tables look as in table 3

The candidates challenge the election authority to open a random half of the ballots, say the ones numbered 2, 4 and 5. The EA reveals the requested information, and the tables look as in table 4. Ballots 2,4 and 5 cannot be used in the election and are excluded from any further representation of the tables (see table 7).

Assume that the choices of the voters are as follows. On ballot 1, the leftmost mark was marked, and the top page was chosen. On ballot 3, the rightmost mark and the bottom page were chosen, and on ballot 6, the leftmost mark and the top page were chosen. Because the canonical ballot is “ab”, “ab” (that is, “ab” on top and bottom pages), left is associated with “a”, and right with “b”. The voter choices eventually end up in  $P_3$ , and when they do, each row describes what can be learned through knowledge of the ballot page chosen by the voter.

The EA performs the first flip to ballots 1,3 and 6 to obtain the partially decrypted ballots as in  $D_3$ , and the totally decrypted ballots as in  $R_1$  (see table 6). The ballots in both  $D$  and  $R$  are shuffled, so it is not possible to link rows among Tables  $P$ ,  $R$  and  $D$ . Post election, the auditor asks the EA to open either the left or the right side of  $D$  (but not both). If the election authority cheats, the auditor will catch it with probability 0.5 (for a higher probability see section ??). In our example, suppose the auditor chooses the right hand side. The EA then reveals  $D_4$  and  $D_5$ . The auditor can now check that  $D_3 + D_4 = R_1$ , and that the commitment  $CD_{4,5}$  to the columns  $D_4$  and  $D_5$  is valid.

## 6 A more technical description

This section provides a more technical description of PunchScan.

Ballot ID	$P_1$	$P_2$	$P_3$	$CP_1$	$CP_2$
1	ab	ab		$C_{1,1}$	$C_{1,2}$
2	ab	ba		$C_{2,1}$	$C_{2,2}$
3	ba	ab		$C_{3,1}$	$C_{3,2}$
4	ba	ba		$C_{4,1}$	$C_{4,2}$
5	ab	ba		$C_{5,1}$	$C_{5,2}$
6	ba	ab		$C_{6,1}$	$C_{6,2}$

$D_1$	$D_2$	$D_3$	$D_4$	$D_5$	$DC$
6	$\rightarrow$		$\odot$	5	$C_A$
5	$\odot$		$\rightarrow$	4	$C_B$
2	$\odot$		$\rightarrow$	1	$C_C$
1	$\odot$		$\odot$	3	$C_D$
4	$\rightarrow$		$\rightarrow$	2	$C_E$
3	$\rightarrow$		$\odot$	6	$C_F$
$CD_{1,2}$			$CD_{4,5}$		

$R_{id}$	$R_1$
1	
2	
3	
4	
5	
6	

Table 2: *PDR* tables as the Election Authority sees them, with all the information available. The tables are properly formed, because, for all the ballots,  $D_2 + D_4$  correctly represents whether  $P_2$  is a flipped version of  $P_1$  or not. For example, for ballot number 3, on the top page, “a” is associated with “Yes”, and b with “No”. On the bottom page, the order is “ba”, thus  $P_2$  is a flipped version of  $P_1$ . In the  $D$  table, in the row corresponding to 3, we have  $\rightarrow + \odot = \text{flip}$ . For ballot 1,  $C_{1,1}$  is a commitment to  $P_1$ ,  $C_{1,2}$  is a commitment to  $P_2$  and so on.

Ballot ID	$P_1$	$P_2$	$P_3$	$CP_1$	$CP_2$
1				$C_{1,1}$	$C_{1,2}$
2				$C_{2,1}$	$C_{2,2}$
3				$C_{3,1}$	$C_{3,2}$
4				$C_{4,1}$	$C_{4,2}$
5				$C_{5,1}$	$C_{5,2}$
6				$C_{6,1}$	$C_{6,2}$

$D_1$	$D_2$	$D_3$	$D_4$	$D_5$	$DC$
					$C_A$
					$C_B$
					$C_C$
					$C_D$
					$C_E$
					$C_F$
$CD_{1,2}$			$CD_{4,5}$		

Table 3: *PD* tables in the Preelection phase, as the public sees them.

Ballot ID	$P_1$	$P_2$	$P_3$	$CP_1$	$CP_2$
1				$C_{1,1}$	$C_{1,2}$
2	ab	ba		$C_{2,1}$	$C_{2,2}$
3				$C_{3,1}$	$C_{3,2}$
4	ba	ba		$C_{4,1}$	$C_{4,2}$
5	ab	ba		$C_{5,1}$	$C_{5,2}$
6				$C_{6,1}$	$C_{6,2}$

$D_1$	$D_2$	$D_3$	$D_4$	$D_5$	$DC$
					$C_A$
5	$\odot$		$\rightarrow$	4	$C_B$
2	$\odot$		$\rightarrow$	1	$C_C$
					$C_D$
4	$\rightarrow$		$\rightarrow$	2	$C_E$
					$C_F$
$CD_{1,2}$			$CD_{4,5}$		

Table 4: *PD* tables after the election authority has replied to the request to open ballots 2,4 and 5

Ballot ID	$P_1$	$P_2$	$P_3$
1			
3			
6			

$D_1$	$D_2$	$D_3$	$D_4$	$D_5$
$CD_{1,2}$			$CD_{4,5}$	

Table 5: Ballots that can be used by voters in the election day. The other ballots were spoiled during the pre election phase. The row commitments are not shown anymore because they won’t be checked, since no other complete row will ever be opened.

Ballot ID	$P_1$	$P_2$	$P_3$
1	ab		a
3		ab	b
6	ba		a

$D_1$	$D_2$	$D_3$	$D_4$	$D_5$
		a		
		b		
		b		
$CD_{1,2}$			$CD_{4,5}$	

$R_{id}$	$R_1$
3	a
5	b
6	a

Table 6: *PDR* snapshot after the polls close. One cannot say what row in the  $D$  table corresponds to what row in the  $P$  or  $R$  table, because the rows are permuted. Thus, the secret ballot principle is satisfied.

Ballot ID	$P_1$	$P_2$	$P_3$
1	ab		a
3		ab	b
6	ba		a

$D_1$	$D_2$	$D_3$	$D_4$	$D_5$
		a	⊙	5
		b	⊙	3
		b	⊙	6
$CD_{1,2}$			$CD_{4,5}$	

$R_{id}$	$R_1$
3	a
5	b
6	a

Table 7: *PDR* snapshot after the post election audit. The election authority was asked to open the right hand side of the  $D$  table. Anyone can check that the intermediary result transformed by  $D_4$  gives the result in  $R_4$  ( $D_3 + D_4 = R$ ), thus the election authority did not cheat. Also  $CD_{4,5}$ , the commitment to  $D_4$  and  $D_5$  is checked. Note that there is still no link between  $P$  and  $R$ , thus the privacy

## 6.1 The ballot

Let  $S$  be a set of symbols. The symbols in  $S$  will appear on both the top and bottom page. We assume that  $S$  is sorted and the order is fixed. We denote by “canonical ballot” a ballot that will have  $S$  (ordered) on both the top and bottom page. Let  $T_p$  (top permutation),  $B_p$  (bottom permutation), and  $D_2$  be three random, independent permutations of  $S$  (in an implementation, the permutation would be pseudorandomly generated as described in section 8).

Compute  $D_4$  such that  $T_p \circ B_p \circ D_2 \circ D_4 = I$  (the composition of the four permutations is the identity permutation). Therefore,  $D_4 = D_2^{-1} \circ B_p^{-1} \circ T_p^{-1}$ .

## 6.2 The tables

We describe the *PDR* tables using notation from relational algebra. In databases relational algebra is heavily used. It has the notions of relations (tables), projections ( $\pi$  - SQL SELECT), selection ( $\sigma$  - SQL WHERE) and join ( $\bowtie$ ). In a relation  $R(A, B)$ ,  $A \rightarrow B$  means that  $A$  implies  $B$  (given  $A$ ,  $B$  is uniquely identified).  $A$  is called a key of relation  $R$ .

Let  $P$  (print) be the following relation:

$$P(B_{id}, P_1, P_2, P_3, CP_1, CP_2), B_{id} \rightarrow (P_1, P_2, P_3, CP_1, CP_2)$$

where  $B_{id}$  is the ballot id (the serial number of the ballot),  $P_1$  is  $T_p$ ,  $P_2$  is  $B_p$ ,  $P_3$  is a projection of  $T_p \circ B_p$  (voter choices),  $CP_1$  is a commitment to  $P_1$ , and  $CP_2$  is a commitment to  $P_2$ . The commitments are cryptographic commitments (see section 9.2 for details).  $P$  contains  $2n$  records.

Let  $D$  (decrypt) be the following relation:

$$D(D_1, D_2, D_3, D_4, D_5, DC), D_1 \rightarrow (D_2, D_3, D_4, D_5, DC)$$

where  $D_1$  is a foreign key pointing to the  $B_{id}$  attribute of  $P$ ,  $D_5$  is a foreign key pointing to the  $R_{id}$  attribute of  $R$  (see below),  $D_2$  and  $D_4$  are permutations of  $S$  described above,  $D_3$  is  $P_3 \circ D_2$ , and  $DC$  is a commitment to the tuple  $(D_1, D_2, D_4, D_5)$ .  $D$  contains  $2n$  records.

Let  $CD$  (commitments to the columns of  $D$ ) be the following relation:

$$CD(CD_{1,2}, CD_{3,4})$$



This relation has only one record.  $CD_{1,2}$  is a commitment to  $D_1$  and  $D_2$ ;  $CD_{4,5}$  is a commitments to  $D_4$  and  $D_5$ .

Let  $R$  (results) be the following relation:

$$R(R_{id}, R_1), R_{id} \rightarrow (R_1)$$

where  $R_{id}$  is a unique identifier and  $R_1$  is  $P_3 o D_2 o D_4$ .  $R$  contains  $2n$  records.

To select all the information for a ballot, we do:

$$(P \bowtie_{B_{id}=D_1} D) \bowtie_{D_5=R_{id}} R$$

### 6.3 The time line

Before the election the election authority(EA) computes  $P(B_{id}, P_1, P_2, CP_1, CP_2)$ ,  $D(D_1, D_2, D_4, D_5, DC)$ ,  $CD(CD_{1,2}, CD_{4,5})$  and makes public  $P(B_{id}, CP_1, CP_2)$ ,  $D(DC)$  and  $CD(CD_{1,2}, CD_{4,5})$ .

In the preelection audit, the auditor randomly selects half of the records in  $P$ . The election authority (EA) reveals  $P \bowtie_{B_{id}=D_1} D$  for all the requested records. The auditor can check that  $P_1 o P_2 o D_2 o D_4 = S$  and that the commitments  $CP_1$ ,  $CP_2$ , and  $DC$  are valid.

During the election, the voters fill in  $P_3$ .

After the election, EA computes  $D_3 = P_3 o D_2$  and  $R_1 = D_3 o D_4$  and makes  $D_3$  and  $R_1$  public.

In the post election audit, the auditor asks EA to either reveal  $(D_1, D_2)$  or  $(D_4, D_5)$ , but not both. EA reveals the requested information. The auditor can either check that  $P_3 o D_2 = D_3$  (using  $P \bowtie_{B_{id}=D_1} D$ ) or  $D_3 o D_4 = R_1$  (using  $D \bowtie_{D_5=R_{id}} R$ ). The chance of EA cheating and not being caught is 50% (see section 6.4).  $CD_{1,2}$  and  $CD_{4,5}$  are also checked.

### 6.4 Multiple instances of $D$

Because EA can cheat and not get caught with 50% probability, we introduce multiple instances of  $D$ . Thus we modify the relation  $D$  as follows: Let  $D$  (decrypt) be the following relation:

$$D(i, D_1, D_2, D_3, D_4, D_5, DC), (i, D_1) \rightarrow (D_2, D_3, D_4, D_5, DC)$$

where  $i$  is the instance number and the rest are as described in section 6.2

Let  $CD$  (commitments to the columns of  $D$ ) be the following relation:

$$CD(i, CD_{1,2}, CD_{3,4}), i \rightarrow (CD_{1,2}, CD_{3,4})$$

where  $i$  is a foreign key pointing to the  $i$  attribute of  $D$ .

In the post election audit, we can now make  $k$  challenges, where  $k$  is the number of  $D$  instances. The auditor will ask to open either  $(D_1, D_2)$  or  $(D_4, D_5)$  for each instance of  $D$ . The chance that EA cheats and does not get caught is one out of  $2^k$ . Thus we can make it as low as we want by increasing  $k$ .

### 6.5 Multiple-question Ballots

We have been implicitly assuming that there is only one question on the ballot per election. The situation becomes slightly more complicated if this is not the case. Punchscan works just fine for multiple-question ballots but the decrypted ballots will preserve the ‘‘cross-question’’ relationships: for example, if 90% of the people who voted for Alice for Governor also voted for Bob for President, the results will reflect this. PunchScan can be extended to hide these correlations if desired.

Trivially, of course, if Punchscan works for one-question elections then we can conduct an  $n$ -question election by giving each voter  $n$  one-question ballots. If we want to preserve the cross-question relationship among two or more questions (perhaps if someone voted ‘No’ for a recall election they are not allowed to vote for a replacement candidate) then we could move those questions onto the same ballot. This would work but seems to us to be not as good (from a ballot design, system overhead, and printing cost point of

view) as the case when we are using one ballot and running one election.

However, we can readily modify this scheme to fix that problem. Suppose we are running  $n$  one-question elections. That is, each voter receives one ballot for each of  $n$  elections and votes, then the votes are counted separately for each election. In this situation, there is one  $P$ -table and one set of  $D$ -tables (and associated  $R$ -table) for each of the  $n$  elections. Let us note that there is no information contained in the  $D$ -tables for election  $A$  that can be used to decrypt the ballots for election  $B$ . Since the shuffles for each election are also independent, we do not need to obscure the link between voter  $x$ 's encrypted ballots in election  $A$  and  $B$ , because when they are decrypted the shuffling will obscure the cross-question relationship for us. In other words, we can print the contests of these ballots together, on the same piece of paper, with the same serial number (and the same  $P$ -table row), just as in the original scheme that reveals the correlations. Because the ballots are decrypted separately, this does not provide any more information regarding the cross-question relationships.

## 7 Proofs

This section contains proofs of some security properties of PunchScan.

### 7.1 Privacy

In this context, the maintenance of privacy requires that an observer's probability distribution of the contents of a given ballot  $i$  (i.e.: the value of voter  $i$ 's vote) be unchanged by observation of the cryptographically-hidden data. In other words,

$$p(b_i|PDR) = p(b_i|R),$$

where  $b_i$  is the value of ballot  $i$ ,  $PDR$  is the entire publicly-observable ballot data matrix, and  $R$  is just the results column of that matrix.

#### 7.1.1 Attacks on $P$

The most straightforward way for an attacker to use the secret parts of  $PDR$  to reveal the vote of voter  $i$  would be to simply decrypt  $P_{1,i}$  and  $P_{2,i}$  and use those to decode  $P_{3,i}$ . If the attacker is unable to break this cryptography, then learning  $P$  would not affect his probability distribution on  $b_i$ . This cryptography can be made as strong as is necessary to protect privacy to any desired level of computational security.

#### 7.1.2 Attacks on $D$

Another method would involve an attack on the shuffle; that is, decrypting the unrevealed link between  $P$  and  $D$  ( $D_1$ ) or between  $D$  and  $R$  ( $D_5$ ). However, the same cryptography is used to secure those columns of  $D$ , so again, an attacker unable to break the cryptography could not learn anything useful from  $D$ .

### 7.2 Integrity

There are four elements of the PunchScan process that are vulnerable to some extent to manipulation of the vote tally by the Election Authority.

- The ballots may be improperly formed.
- The ballots may be improperly printed.
- The ballot markings may be improperly recorded.
- The marked ballots may be improperly decrypted.

Each of these vulnerabilities is addressed by an audit procedure.

### 7.3 First Audit

The first audit procedure ensures that the ballots are well-formed, meaning that for each ballot,  $P_1 + P_2 = D_2 + D_4$  for the row in each  $D$ -matrix associated with that ballot. This involves spoiling some fraction of the ballots by unlocking this secret data.

In general, suppose there are  $n$  ballots, the election authority has cheated by malforming  $k$  of them, and  $f$  ballots are chosen at random to be examined. The probability that the election authority gets away with this cheat is:

$$\frac{\text{thenumberofpossibilitieswheretheauditorchoosesonlyvalidvotes}}{\text{allthepossiblechoices}}.$$

The number of all the possible choices is  $\binom{f}{n}$  ( $n$  choose  $f$ ). The number of ways to choose  $f$  all valid votes from a total of  $n$  where  $k$  of them are invalid, is  $\binom{f}{n-k}$  (choose  $f$  votes out of  $n-k$  that are valid). So

$$\text{the election authority cheats and gets away with it with a probability } p = \frac{\binom{f}{n-k}}{\binom{f}{n}} = \frac{\frac{(n-k)!}{f!(n-k-f)!}}{\frac{n!}{f!(n-f)!}} = \frac{\frac{(n-k)!}{(n-k-f)!}}{\frac{n!}{(n-f)!}}.$$

Note that  $f + k < n$ , so that  $n - k - f > 0$  and  $(n - k - f)!$  exists and it isn't in the special case of  $0!$ . If  $f + k > n$  the probability is 0.

In the interest of simplicity, from here we may compute two upper bounds on the chance that this attack will not be detected:

1.  $\frac{(n-k)!}{n!} \times \frac{(n-f)!}{(n-k-f)!} = \frac{1}{n \times (n-1) \times \dots \times (n-k+1)} \times (n-f) \times (n-f-1) \times \dots \times (n-f-k+1) = \frac{(n-f) \times (n-f-1) \times \dots \times (n-f-k+1)}{n \times (n-1) \times \dots \times (n-k+1)} = \frac{n-f}{n} \times \frac{n-f-1}{n-1} \times \dots \times \frac{n-f-k+1}{n-k+1} = (1 - \frac{f}{n}) \times (1 - \frac{f}{n-1}) \times \dots \times (1 - \frac{f}{n-k+1}) < (1 - \frac{f}{n})^k$
2.  $\frac{(n-k)!}{(n-k-f)!} \times \frac{(n-f)!}{n!} = (n-k) \times (n-k-1) \times \dots \times (n-k-f+1) \times \frac{1}{n \times (n-1) \times \dots \times (n-f+1)} = \frac{(n-k) \times (n-k-1) \times \dots \times (n-k-f+1)}{n \times (n-1) \times \dots \times (n-f+1)} = \frac{n-k}{n} \times \frac{n-k-1}{n-1} \times \dots \times \frac{n-k-f+1}{n-f+1} = (1 - \frac{k}{n}) \times (1 - \frac{k}{n-1}) \times \dots \times (1 - \frac{k}{n-f+1}) < (1 - \frac{k}{n})^f$

Thus, our upper bound on the probability that the Election Authority gets away with malforming  $k$  out of  $n$  ballots when  $f$  of those ballots are audited is  $\min[(1 - \frac{f}{n})^k, (1 - \frac{k}{n})^f]$ .

### 7.4 Second Audit

In order to check that a given ballot receipt was properly printed, one can reencrypt it (that is, recompute the commitments) and compare it with the  $P$ -matrix. Suppose  $n$  ballots remain unspoiled after the First Audit,  $f$  are actually used by voters who later check the commitments, and  $k$  of them are improperly printed. Once again, the upper bound on the probability that none of the misprinted ballots are detected is  $\min[(1 - \frac{f}{n})^k, (1 - \frac{k}{n})^f]$ .

### 7.5 Third Audit

In addition to checking that the ballot is correctly printed, one can also check whether the recorded ballot mark matches the mark on the receipt. In effect, this verifies the correctness of  $P_3$ . Again, if  $n$  ballots remain unspoiled after the First Audit,  $f$  are actually used by voters who verify that their ballot marks are correctly recorded online, and  $k$  ballots are incorrectly recorded, then the upper bound on the probability that none of the incorrectly-recorded marks are detected is  $\min[(1 - \frac{f}{n})^k, (1 - \frac{k}{n})^f]$ .

### 7.6 Fourth Audit

The Election Authority may influence the vote tally by incorrectly decrypting the ballots. There are two methods we may use for auditing the Election Authority to ensure that this does not occur.

### 7.6.1 Ballot-wise Auditing

Suppose the auditor goes through a  $D$ -matrix ballot-by-ballot (that is, row-by-row) and randomly chooses whether to inspect (open) the “left” or “right” commitment for each ballot. This situation is different from the first three audits because all ballots are inspected, but each inspection has only a  $\frac{1}{2}$  chance of catching a modification. This makes the situation simpler; the chance of  $k$  modified ballots all escaping detection is  $\frac{1}{2^k}$ .

### 7.6.2 Table-wise Auditing

On the other hand, the auditor may choose to open all the “left” or “right” commitments for a given  $D$ -matrix. Assuming that the Election Authority intends to cheat during the decryption and is aware of this, he will put all his cheating in a given  $D$ -matrix in either the “left” or “right” commitment, so that he has a  $\frac{1}{2}$  chance of escaping detection when that  $D$ -matrix is inspected. If there are  $n$   $D$ -matrices, then the chance of escaping detection if any ballots are incorrectly decrypted is  $\frac{1}{2^n}$ .

### 7.6.3 Comparison

Both methods of auditing have desirable properties. The Ballot-wise method has the feature that the probability of detecting cheating is a function of the number of ballots cheated on, and increases exponentially with a linear increase in number of cheated ballots. The Table-wise method has the feature that the audit does not reduce the size anonymity set created by the shuffle.

## 8 Permutations

PunchScan requires two types of permutations to be generated:

- row permutations
- mark permutations

Row permutations refer to the permutations of the rows of the  $D$  table and mark permutation refer to the order in which the positions are associated with marks on the ballot and to  $D_2$  and  $D_4$ .

### 8.1 Row Permutations Generation

Consider an “unshuffled”  $D$ -matrix where  $D_1 = [1, 2, \dots, 2n]$ , so row  $x$  of  $PDR$  represents ballot  $x$  across the entire row, and  $D_5$  is blank. The EA should generate this matrix as the first step; call it  $\delta$ . Generating the row permutations will therefore take the form of the generation of  $D^1 \dots D^{n_D}$ , where  $D^i$  denotes the  $i^{th}$  shuffled  $D$ -matrix.

The  $D$ -matrices will be generated from  $\delta$  as follows:

1. Randomly shuffle the rows of  $\delta$ ; call this  $D^1$ .
2. Let  $D_5^1$  equal a random shuffling of  $\{1, 2, \dots, 2n\}$ .
3. For each  $i$  from 2, 3,  $\dots$ ,  $2n$ , let  $D^i$  equal a random shuffling of the rows of  $D^1$ .

This involves  $n_D + 1$  permutations of  $\{1, 2, \dots, 2n\}$ . It should be clear that if  $(y, D_1^i) = x$  and  $(y, D_5^i) = z$ , then for all  $j$ ,  $(y, D_1^j = x)$  implies that  $(y, D_5^j) = z$ ; in other words, since each row of  $D^1$  contains a pointer to a (unique) row (ballot) of  $P$  in  $D_1$  and a (unique) pointer to  $R$  in  $D_5$ , reordering its rows does not change the destination (in  $R$ ) of any ballot in  $P$ .

## 8.2 Implementation

### 8.2.1 Permutation Algorithm

We will use the following permutation algorithm to permute the unshuffled matrix. This algorithm will generate a permutation  $\pi$  of  $1, 2, \dots, m$ , given as input  $m$ , some encryption function  $E$ , and some key  $K$ .

First, create a table with  $m$  rows and 2 columns. Populate column 1 of the table with  $1, 2, \dots, m$  and column 2 of the table with  $E_K(1), E_K(2), \dots, E_K(m)$ ; in other words,  $(i, 2) = E_K((i, 1))$  for every row  $i$ . Next, sort the table according to column 2. Let  $\pi(i) = (i, 1)$ ; column 1 is now a permutation of  $1, 2, \dots, m$ .

If the key  $K$  were generated randomly, and the function  $E$  is a good encryption algorithm, then the permutation output by the algorithm will be random. (That is, it will preserve any randomness in  $K$ .)

### 8.2.2 Application of the Algorithm

The EA can use this algorithm to implement the  $D$ -matrix generation algorithm above as follows:

1. Generate a permutation  $\pi_{D^1}$  of  $1 \dots 2n$ . Let  $D_1^1 = \delta$ , sorted by  $\pi_{D^1}$ ; that is, row  $x$  of  $\delta$  becomes row  $\pi_{D^1}(x)$  of  $D^1$ .
2. Generate a permutation  $\pi_R$  of  $1 \dots 2n$ . Let  $D_5^1 = \pi_R$ .
3. For each  $i$  from 2 to  $n_D$ , create  $D^i$  by generating a permutation  $\pi_{D^i}$  of  $1 \dots 2n$ . Let row  $y$  of  $D^1$  become row  $\pi_{D^i}(y)$  of  $D^i$ .

## 8.3 Mark Permutations

The mark permutations, in contrast, are much simpler to generate. In order to produce all possible associations of candidate names with ballot symbols, it is not necessary to randomly permute both lists; it is only necessary to cyclically shift both lists a (different) random amount. So, to generate the mark permutations for ballot  $x$ , where the ballot has  $c$  candidate names on the top page and  $c$  mark symbols on the bottom page, the EA need only generate two random numbers between 1 and  $c$ , and record these numbers as  $P_1$  and  $P_2$  to indicate the shift distance for the pages of ballot  $x$ .

Each  $D$ -matrix instance will require its own set of decrypting mark permutations (columns  $D_2$  and  $D_4$ ). (It is for this reason that at least the decrypting mark permutations must be performed after the row permutations.) For each row of each  $D^i$ , the EA generates a random number between 1 and  $c$ , and records this number in  $D_2^i$ .  $D_4^i$  is set such that the modular sum of the ballot's entries in  $P_1$  and  $P_2$  equals the sum of its entries in  $D_2$  and  $D_4$ .

### 8.3.1 Random Number Generation

The permutation algorithm described above can also be used for the random number generation. The Election Authority can compute a permutation  $\pi$  of  $1, 2, \dots, c$  and use  $\pi(1)$  as the random number.

## 9 Commitments

This section describes how the commitments in PunchScan are computed. Comma (“,”) stands for concatenation. There are two AES 128-bit keys secret  $MK_1$  and  $MK_2$ , and a public 128-bit constant,  $C$ .

### 9.1 Computing AES keys

This section requires the use of two 128-bit AES keys. Given message  $M$ , let  $M_{128}$  be the first 128 bits of  $M$  (if  $M$  is shorter than 128 bits,  $M$  will be padded with trailing zeros) a random key  $SK_m$  is generated as follows:

$$SK_m = D_{MK_1}(C \oplus E_{MK_2}(C \oplus E_{MK_1}(M_{128})))$$

where  $\oplus$  is the XOR operation and  $E$  and  $D$  are AES Encrypt and Decrypt EBC NoPadding operations.

## 9.2 Commitment Algorithm

Given a message  $M$ , the commitment to  $M$  is computed as follows:

1. generate a 128-bit AES key  $K_m$  as described in 9.1
2. encrypt the public constant  $C$  with  $K_m$ , using AES 128-bit ECB NoPadding. Let the result be  $SK_m = AES_{K_m}(C)$ .  $SK_m$  has 128 bits.
3. concatenate  $M$  with  $SK_m$  and hash everything using SHA256, resulting in  $h_1$ . So,  $h_1 = SHA256(M, SK_m)$ ;
4. let  $h_2 = SHA256(M, AES_{SK_m}(h_1))$ , where the AES encryption is AES 128bit ECB PKCS#5Padding
5. the commitment is  $h_1, h_2$  ( $h_1$  concatenated with  $h_2$ )

Below is how to compute  $M$  for all the commitments needed in PunchScan.

### 9.2.1 $M$ for $P_1$

$M$  is obtained by concatenating the serial number of the ballot to a constant particular to  $P_1$  and with the text on the top page of the ballot.

$M = i, "P1", P_1$  where  $i$  is a string representing the serial number of the ballot, " $P1$ " is a constant string (capital P concatenated with digit 1) and  $P_1$  is the string in  $P_1$  (the string representation of the top page)

### 9.2.2 $M$ for $P_2$

$M$  is obtained by concatenating the serial number of the ballot to a constant particular for  $P_2$  and with the text on the bottom page of the ballot.

$M = i, "P2", P_2$  where  $i$  is a string representing the serial number of the ballot, " $P2$ " is a constant string (capital P concatenated with digit 2) and  $P_2$  is the string in  $P_2$  (the string representation of the bottom page)

### 9.2.3 $M$ for rows in $D$

$M$  is obtained by concatenating all the known values in a row in  $D$ . The known values are: the pointer to the  $P$  table ( $D_1$ ), the first mark permutation ( $D_2$ ), the second mark permutation ( $D_4$ ) and the link to the  $R$  table ( $D_5$ ).

$M = D_1, D_2, D_4, D_5$  all being string representation of fields in  $D$

### 9.2.4 $M$ for columns in $D$

$M$  is obtained by concatenating all the values in the first column and then concatenating all the values in the second column.

For the leftmost columns

$M = D_{1,1}, D_{2,1}, D_{3,1}, \dots, D_{n,1}, D_{1,2}, D_{2,2}, D_{2,3} \dots D_{n,2}$  all being string representations

For the right most columns

$M = D_{1,4}, D_{2,4}, D_{3,4}, \dots, D_{n,4}, D_{1,5}, D_{2,5}, D_{2,5} \dots D_{n,5}$  all being string representations

We only need to protect two 128-bit AES keys,  $MK_1$  and  $MK_2$  in order to preserve the security of the system.

Note that the public cannot verify that the AES keys have been generated in this way, or rather in some other way. Therefore, this system unfortunately introduces a potential covert channel via the AES keys.

## 10 Acknowledgments

We would like to thank David Chaum, Poorvi Vora, Rick Carback, Jeremy Robin and Ben Adida, for the vibrant discussions and insightful comments.