usenix ;login:
DECEMBER 2012 VOL. 37, NO. 6
ELECTRONIC SUPPLEMENT

Conference Reports

2012 Electronic Voting Technology Workshop/ Workshop on Trustworthy Elections (EVT/ WOTE '12)

Bellevue, WA
August 6–7, 2012

New Interfaces

Summarized by Harvie Branscomb (harvie@electionquality.com)

Operator-Assisted Tabulation of Optical Scan Ballots

Kai Wang, University of California, San Diego; Nicholas Carlini, Eric Kim, Ivan Motyashov, Daniel Nguyen, and David Wagner, University of California, Berkeley

Kai Wang discussed OpenCount (code.google.com/p/opencount), an open-source software project to productively combine human cognition with machine functionality for the purpose of tabulating scans of paper ballots. The design was motivated by a need to go beyond what software can do to interpret interesting cases, such as poor markings and erasure marks. OpenCount interleaves computer vision techniques with focused operator verification to produce a "cast vote record" of each scan of a ballot suitable for performing "single-ballot level" risk-limiting comparison audits. The system does not rely on ballot vendor specifications or definition files and may be used with existing scans of voted ballots, but it requires a not-voted or "blank" instance of every unique ballot style used in the election for configuration.

The first phase involves an interaction with an operator to identify a rectangular area around a voting "target" that the system uses to find similar others, automatically grouping them into clusters. Portions of text or image are operator selected to identify and classify each unique style, such as party, language, precinct, etc. In the second phase, voted ballot scans are spatially translated and rotated as necessary for registration with the data from the unvoted examples at the pixel level across the entire ballot, and then again for each voting target. Each target in every voted ballot is displayed in an array ordered by average pixel density to allow the operator to inspect visually and determine the threshold between marks to be classified as votes and those not to be so classified.

OpenCount has been successfully validated in California counties through secondary scanning of several manufacturers' styles of ballots in five risk-limiting audits in 2011 and two in 2012, with at least four more upcoming. Election officials have agreed that OpenCount provides a more accurate count than purely machine counts not using operator input.

Anna Queredo asked how the system notifies the operator which ballots to look at. Kai explained that the operator scrolls to the border between marked and unmarked targets and focuses attention there. He said there is also a function to handle ballot scans for which something unusual happened separately. Jeremy Epstein asked how to notice marks that are outside the "target." Kai explained that marks within a defined rectangular area surrounding the "target" are recognized, but everything outside of that rectangle is ignored.

A Hybrid Touch Interface for Prêt à Voter

Chris Culnane, University of Surrey, Trustworthy Voting Systems Project

Prêt à Voter (pret ah votay) in its original form is an end-to-end verifiable paper-ballot voting system design that is machine tabulated such that no machine learns what the voter intent is; thus, it systematically retains the privacy of the vote. The design's central element is a paper ballot that can be split so that, after the ballot is marked, the randomly ordered list of candidates becomes separated from the voter's marks. A crypto key containing a signed serial number protects access to the knowledge of the order of the candidates on the ballot while the marks themselves remain public. Chris Culnane's talk introduced an accessibility extension of Prêt à Voter in which the right-hand side of the ballot, the portion to be marked, is implemented on a touch screen such that all of the integrity features of the original design are maintained while additional accessibility features such as tactile and auditory cues could be implemented.

Two implementations were described, one for the original Microsoft Surface interactive desktop and another for a 3M Multi-Touch M2256PW. Chris also speculated about a third using the Samsung SUR40 with Microsoft PixelSense. In the design, the surface of the screen must recognize a 2D barcode or a coded conductive ink or foil that (1) informs the system of the location and orientation of one or more paper left-hand ballot sides containing lists of human-readable ballot choices (e.g., candidates) and (2) allows the screen to display the right-hand side(s) of the ballot as indistinguishable vote targets in the appropriate location(s). In Chris' implementation only the left-hand side of the Prêt à Voter ballot exists in paper form.

Chris admitted to concerns that voters might believe the system could recognize their face through the glass, although the technology does not have that capability. Philip Stark asked about the time-frame to deploy. Chris admitted that he has limited access to the necessary equipment and much work is yet to be done. Jeremy Epstein asked about the range of disabilities that could be served by this system. Chris said the system only requires the ability to place the left-hand side of a Prêt à Voter paper ballot on the display surface and read

it; thus embossing such as Braille may be needed. Also, the system must hold the paper in place, so this limits the extent to which the display can be placed vertically. Flexibility in orientation is advantageous to voters with disabilities. Chris suggested that a move to a smaller form factor would help broaden the scope of application to various disabilities, and price is also a concern. Peter Neumann questioned the need for trust of the underlying technology, but Chris reassured him that the system does not learn the permutations of the candidate order; instead the system has access to only a serial number protected by cryptography.

Election Auditing

Summarized by Harvie Branscomb (harvie@electionquality.com)

A Bayesian Method for Auditing Elections

Ronald L. Rivest and Emily Shen, Massachusetts Institute of Technology

Ron Rivest brings new resources from the field of statistics to the practice of auditing elections. His talk about a "ballot-polling" method of auditing described the use of Bayesian methods and their multiple advantages, including for more familiar "comparison audits." Ballot polling does not require access to data from a voting system and instead independently predicts the likelihood that any given candidate would be declared the winner after counting all of the ballots while usually counting relatively few. It does require the ability to randomly select and interpret the voter intent on every ballot marked by every voter in an election contest. While the method is easy to describe on a single page, the extent of the calculations needed requires the assistance of a machine for most elections. The Bayes audit does not require knowledge of the margin of victory and conveniently permits multiple auditors with multiple "Bayes priors" to be accommodated. Bayes priors can reflect real biases among interested parties, such as the expectations of a losing candidate who believes that uncounted ballots are voted in his or her favor.

Ron reported that Bayes audits offer good efficiency, comparable to that found in Stark's ballot-polling and single-ballot comparison-audit methods. He said Bayes methods can also be applied to comparison audits, which offer even better efficiency over ballot polling. Many voting methods can be supported. Small and controllable miscertification rates are observed. Even if the audit is stopped early for practical reasons, meaningful results can be obtained. Disadvantages include applicability only to single-ballot audits with results depending somewhat on the choice of prior. How Bayes audits relate to risk-limiting audits remains an open question.

David Flater was concerned about non-obvious stopping criteria and the need to control the risk to a specified level. Ron Rivest explained that the analysis involved in the Bayes audit is nicer than with other methods where it gets complicated but that risk measures depend on the priors. Bayes methods represent a solid approach to, for example, financial audits.

Peter Neumann expressed concern that this method has the potential to predict election outcomes from incomplete data and also that IRV outcomes could not be calculated without complete data. Ron Rivest explained that, for all voting methods, you need everything "in" before beginning the audit. John Bodin talked about comparison audits and the need for an identifier to connect ballots to interpretations and the security of this identifier. Philip Stark commented that part of an election is convincing the loser that they have a "prior."

More information is available at:
people.csail.mit.edu/rivest/bayes.

BRAVO: Ballot-polling Risk-limiting Audits to Verify Outcomes

Mark Lindeman, Philip B. Stark, and Vincent S. Yates, University of California, Berkeley

Philip Stark (statistics.berkeley.edu/~stark/Vote) opened by emphasizing the importance of evidence-based elections and the critical need for an adequate evidence trail measured by a compliance audit as a prerequisite to any successful election audit. Stark has been conducting single-ballot risk-limiting audits both in comparison-audit format and the more recently proposed ballot-polling mode. In a risk-limiting audit, the question is not how many ballots to audit at first, but when to stop. If there is compelling evidence that the outcome is correct, then stop; otherwise continue the audit, ballot by ballot, tabulating the result incrementally at each step. In ballot-polling audits, you hand count votes, whereas in comparison audits, you count the discrepancies between a machine count and a hand count. According to Philip, in defining the "risk" of a risk-limiting audit, one assumes that a reported outcome might be wrong in the most maliciously difficult way to detect. The "risk" is the chance that this wrong-outcome scenario would not be detected and would not be corrected by the audit. This is quite different from the risk that any outcome is wrong. Numerous risk-limiting audits have now been conducted by Stark et al. in California elections ranging in size from 200 ballots to 121,000 ballots.

Ballot-polling audits require more ballots to be audited than equivalent comparison audits, but do not require any results from the voting system and have no setup costs such as the need for secondary scanning of ballots. Although polling audits do not check the voting system tabulation, they do expose the voter marks on only relatively few sampled ballots and a number comparable to that of a precinct comparison audit. A good ballot manifest is needed in order to be able to select a random sample, but the method can be executed with dice and a pencil and paper if desired.

The methodology is reminiscent of a public opinion poll where the ballot is asked, "What do you say?" Philip Stark reported several successful election audit experiences in

California. He then extrapolated the audit workload for the average statewide presidential contest. Among 255 statewide presidential contests between 1992 and 2008, the median expected sample size for a statewide ballot-polling audit would be only 307 ballots.

Douglas Wikström suggested a potential for avoiding sequential sampling by using simultaneous multiple ballot sampling or even some special handling of the ballot in the voting booth such as "tossing a p coin." Philip agreed some potential benefit might result. Peter Neumann asked about exit polls. Philip Stark replied that exit polls are a biased sample, encounter problems with people's willingness to answer accurately, and are in effect a mess.