To: Jerome Lovato <jerome.lovato@sos.state.co.us>
From: Ron Rivest <rivest@mit.edu>
Re: Representative Group Meeting 3/3/17 and related matters on
Colorado audits
Date: March 9, 2017

Here are some notes and thoughts in response to the excellent
discussion held 3/3. Much of this is just me trying to record and
organize what is already well understood; thanks to Philip Stark, Neal
McBurnett, Harvie Branscomb, and others for feedback on an earlier
draft.

(1) Key notions:
        County
        Contest
        Collection (of paper ballots)
        Ballot Manifest (for a collection)
        CVRs

    Assumptions:

    (a) A contest may be state-wide, multi-county, county-wide,
        or sub-county in extent. A multi-county contest might be
        sub-county in some of its counties.

    (b) A county will have one or more collections (of paper ballots)
        for ballots cast in that county.

    (c) Each collection has an associated ballot manifest.

    (d) A ballot manifest specifies how many paper ballots are in the
        collection, and where they are located. It also specifies
        the number of ballots in the collection having each contest.

    (e) The ballot manifest should be trustworthy -- not compiled
        using vender-supplied equipment, and not derived from the
        CVRs.

    (f) A collection doesn't have smaller components with the same
        information compiled; otherwise the collection should be divided
        into smaller collections, each with its own ballot manifest.

    (g) A collection may have associated CVRs resulting from the
        scans of its ballots. These give, for each ballot scanned,
        the voter's choice for each contest on the ballot. This
        includes undervote/overvote/etc. information.

    (h) The CVRs are untrustworthy (that is why we are auditing!).
        Not only are the voter's choices untrustworthy, but so are
        the identities of the contests on each ballot, even.

    (i) In the simplest mode, any sampling done will be uniform within
        a collection, and be based on the ballot manifest. We might use

more sophisticated sampling schemes that also use the CVRs for increased efficiency, but these schemes are also a bit more complicated as they need to take into account the possibility of errors in the CVRs regarding the contests on each ballot.

(j) Some collections may have no associated CVRs.

(k) Some jurisdictions may have two-card ballots. It seems simplest to treat each card as a separate ballot with its own ballot style (set of contests). There should be no reason to maintain the connection between two cards submitted by the same voter. With two-card ballots, the ballot manifest should have an entry for each card received in the collection.

(l) The reported election outcomes for each contest are computed from the CVRs (and possibly some other tie-breaking info). The CVRs (and perhaps the tie-breaking info) should all be committed to before the audit starts (when the random sample is selected), so they can't be changed during the audit. The tie-breaking info you may want to commit to before election day closes, especially if more complicated voting methods like IRV are used, when there may be ties not only in the final outcome but also in intermediate steps of determining the final outcome. (This is probably not a concern in Colorado, so breaking ties at the end should be fine.)

(2) Processes:

(a) Ballot reconciliation and accounting:
Produces the ballot manifests.

(b) BM/CVR reconciliation:
Happens before audit starts.
Ensures that there is exactly one CVR per location in the BM.
(Here a ballot location as specified in the BM might contain more than just one card, if dups of cards are attached to the original card, or if somehow the two cards of a two-card ballot scheme are attached, although the latter seems unnecessary.)
(Remove redundant CVRs specifying the same ballot location.)
(Add dummy/zombie ones for locations with no scan data.)
(See Stark https://arxiv.org/abs/1207.3413 "Limiting Risk" for relevant discussion.)

(c) Because the CVRs are untrustworthy, sampling from a collection uniformly using only the ballot manifest as a guide works fine. But this may (or may not) be somewhat inefficient. Other approaches that use the use the CVRs to steer the sampling towards ballots with a particular contest within a collection may be more efficient (if not all ballots have that contest), but they need to account for the possibility that the correspondence between CVRs and ballot manifest locations may be inexact, due to errors in the CVRs, and so a selected ballot might not have the expected contests on it.

(d) A ballot-level comparison audit is the default audit method.

(e) If there is a collection with no CVRs, then a ballot-polling
    method should be used.

(f) If a box of ballots is dropped on the floor and becomes
    disordered, we should just pick them and put them back in
    the box. Sampling within that box will yield errors, since
    the ballots are out of order. But the audit should be able
    to tolerate that, and there won't be many ballots sampled
    from that box anyway.

(g) Similarly, the audit should be able to tolerate the situation
    when we pull the wrong ballot to audit. This will probably
    happen, but at a low rate. Such mistakes will just make the
    audit run longer, since it increases the error rate (rate at
    which CVRs don't match the selected paper ballots). We don't
    need to attempt to detect when we have pulled the wrong
    ballot, unless there is definitive and unique ballot
    identification, such as a printed unique ballot ID on each
    ballot.

(h) When a ballot is selected and examined by hand, the voter's
    choices for all contests should be recorded. (This is a policy
    decision; it could instead say to only record the voter's
    choices for the contests still under audit, but that seems
    more complex and error-prone.)

(i) ALL audit info should be uploaded to the state Audit Center.
    This means all ballot manifests, CVRs, tie-breaking info,
    audit data resulting from hand examination of the paper
    ballots (voter choices on each examined ballot), as well as
    any other adjudications, exceptions, decisions made on
    exceptions, etc. It is important that the ballot manifests
    and CVRs be uploaded to Audit Central before Audit Central
    arranges the sampling seed and sampling details. Sending CVRs
    may be sent as soon as the corresponding collection is
    ``closed'' (no more ballots will be admitted, and no changes
    to the CVRs for that collection of any sort will be made).

(j) I don't know whether for the Colorado audits it is
    best to think of the notions of "county" and "collection"
    as being essentially equivalent (so there is exactly
    one collection per county), or whether there might reasonably
    be more than one collection per county. I suspect the
    latter is true. (And the latter may be desirable if there
    are sub-county audited contests; see (3)(c) below.)
    There may be different collections within a county for
    vote-by-mail, drop-box voting, and in-person voting, for
    example.

(k) Audit Central will guide the audit, by publishing "target
    sample sizes" for each collection (in advance of each day of

the audit), and information about in what order the ballots
for a collection will be audited (which may even allow
counties to ``audit ahead'' a bit).

(l) Contests that are audited "opportunistically" (i.e. not
required to reach specified risk limits but nonetheless having
voter choices observed) will have their risk measured at the
end of the audit.

(3) Other comments:

(a) Ballot images seem irrelevant for the audit. The reported
election outcomes were computed from the CVRs, and the
audit checks the accuracy of the CVRs by hand-comparison
with the corresponding paper ballots. The ballot images
are not part of the critical CVRs-to-outcome path that is
being audited.

That said, the ballot images may have some utility for
PR, for discovering issues different than or beyond
the simple "is the outcome correct?" question the audit
is asking, or for satisfying the public or the candidates
that the audit has credibility.

(b) If a contest is included on ballots in with-CVR collections,
and also on ballots in no-CVR collections, there isn't
great theory for dealing with this (that I know of). One
approach is to just ignore the CVRs available, and do a
ballot-polling audit for this contest. Another is to
associate a "dummy" CVR with each ballot from the no-CVR
collection, and do a ballot-comparison audit. There are
other approaches as well (e.g. some bayesian methods).

(c) When auditing a contest, it is helpful to know if its relevant
collections are "pure" in the sense that they all contain
ballots having the given contest. If the collection is mixed,
with some ballots having the contest and some not having the
contest, the audit may be less efficient. I worry that
contests with small margins will unduly expand the audit in
situations where having smaller more focused collections would
not (since the small-margin contest may be in a smaller
collection). In 2017, only having audits for contests whose
collections of relevant paper ballots are "pure" may be
simplest (which means running risk-measuring audits for any
other contests whose collections aren't "pure"). Sorting of
ballots may be worth considering for elections in 2018 and
beyond...

(d) There are some "privacy concerns" that aren't addressed
in the above remarks.