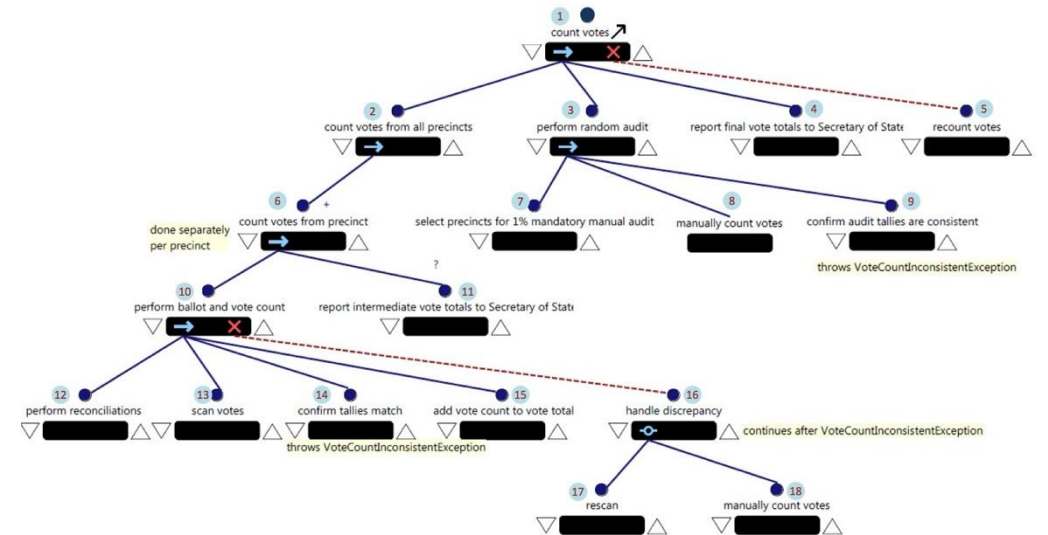


# Lecture 26, June 3, 2026

ECS 235B, Foundations of Computer and Information Security  
Spring Quarter 2026

# What Should an E-Voting System Do?

- Replace manual activity, existing technology used in election process with better technology
  - Better in the sense of improving some aspect of the election process
- Examples
  - Easier to program ballots than print them
  - Can handle multiple languages easily
  - Easier to tally than hand counting



# Assurance

- Provide sufficient evidence of assurance to target audience that using e-voting systems makes elections at least as secure, accurate, etc. as elections without them (that is, using paper ballots)
- Who is “target audience”?
  - Computer scientists, election officials, politicians, *average person*

# Brief History

- Presidential election of 2000: massive confusion over ballots, and counting ballots, in Florida
  - Butterfly ballots did not align properly
  - Hanging chads made determining some votes difficult
- Help America Vote Act appropriated money to pay for electronic voting systems
- Federal standards developed by FEC
  - Voluntary Voting System Guidelines
  - NIST developing next generation

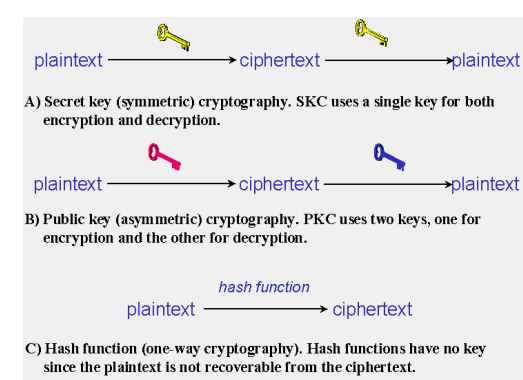
# Problems Developed in Testing

- 2003: Johns Hopkins people analyzed voting system program
- January 2004: RABA study of Diebold systems in Maryland
- April 2004: Diebold made available updates that were not certified
- Summer 2007: CA top-to-bottom review
  - Followed by EVEREST review in Ohio
- 2011: Washington DC internet voting test compromised
  - And the friendly attackers threw out the hostile ones
- 2014: Analysis of Estonia e-voting systems: many vulnerabilities found
- 2020: Voatz mobile voting app based on “blockchain technology”: many vulnerabilities found

# Problems Developed in Use

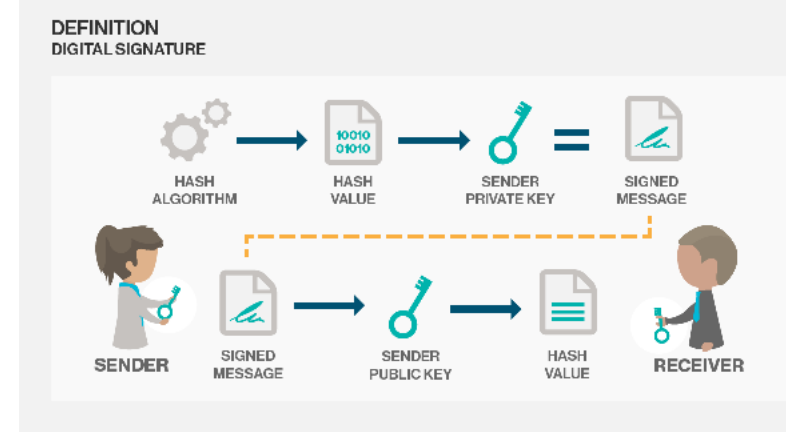
- Boone County, IN, 2003: 144,000 votes cast in a county with about 6,000 voters
- In 2006, polls opened late in several California (CA) counties (San Diego, Alameda, Plumas, Kern, Solano) due to system problems
- December 2006: Florida CD-13 post mortem of massive undervotes in a hotly contested race
- South Bronx, NY, 2010: a scanner miscounted 69/103 (70%) of ballots in Sep., then 156/289 (54%) in Nov.
- Los Angeles, CA, 2020: electronic poll books had connectivity problems, resulting in unacceptably long lines; BMDs failed, had paper jams

# Adding Cryptography



- RABA: Diebold’s implementation of SSL protected confidentiality of precinct results, *but not integrity*
- Yolo County analysis: Hart used “random” access code on eSlates
  - Actually “pseudo-random”, and it took looking at 20 such codes in sequence to regenerate all 10,000 possible codes (same for all systems)

# A Classic Example of Crypto



- Diebold added digital signatures to ballots in the version after the one California reviewed
  - Not examined in TTBR because it wasn't certified in California
- FSU SAIT: Alec Yasinsac and his team examined it
  - Signing technique was flawed, enabling forging of ballots







# When Random Isn't Random

- Hart Intercivic systems have 2 components
  - Hart e-voting system
  - Judge's Booth Controller
- JBC generates a "random" 4 digit number
- Voter goes to e-voting system, enters number, and then can vote
- But numbers are pseudorandom, *not* random
- Students generated 100 numbers, then wrote down the next 100 numbers
  - And verified they were correct

# How to Get There

- You need both standards and testing
- They must be independent of the developers of the systems
- They need to consider the users, operators, and maintainers of the systems
- Reports should show what tested, why, and how
- For e-voting systems, penetration testing is a *must*

# Add in the Internet



- It will enable authorized voters who cannot vote due to distance (or other factors) to do so
- It will increase authorized voter participation
- It will bring our elections into the modern, technological world
- It will be cheaper because we don't have to store the paper ballots

## Problem:

- Election systems are now accessible to many more people than authorized voters!

# Where Would Attackers Strike?

- Probably not regular, individual electronic voting systems
- Attack the voter registration databases to disenfranchise voters
- Or attack the vendors and change the programs that run on those systems, or on the tallying systems



# Internet Elections


- If we can bank over the Internet, why don't we vote?
- Won't it increase election turnout?
- Attack surface increases
  - Election office resources won't increase enough

# And If You Vote via Internet ...

- Is your home PC/Mac secure?
- Is your smartphone secure?
- Are your router, ISP, ... secure?



# And If You Vote via Internet ...



# Remote Voter Verification of Ballots

- Trick here is to protect against the validating mechanism being corrupted
- Example: we examined a system that enabled voters to check that their ballots were recorded correctly, and counted correctly, remotely
  - Used very neat cryptography, done by experts
  - We simply changed the web page on which the information that the user used to do the validation – no cryptography involved!

Moral: attackers don't have to rig or corrupt an election  
They just have to make you *think* they did!

# Blockchains

- Background
  - Take ballot or chain of ballots and compute a hash from them
  - Encrypt this with a cryptographic key you keep secret (private key)
  - Publish the inverse cryptographic key (public key) so others can verify the small value was not changed
- For voting: many proposals for handling the chains

# Why Blockchains Fail for Elections

- Problem #1: denial of service (already discussed)
- Problem #2: how are those cryptographic keys generated?
  - A. Voter generates the pair (this is how it's usually done for other uses), and publishes the public key
  - A'. I vote multiple times, possibly under the name of a different voter each time. Prove I was the one who did this, and determine which votes are mine.
  - A''. I want to sell my vote. I give my private key to the purchaser. She can use the public key to verify that is my private key, and then see how I voted by finding the specific ballot added using that public key.
  - B. Election officials assign key. Now *they* can determine how I voted!

# How Not to Test Voting Over the Internet

- Occasional bills in various legislatures to do a “pilot study” using Internet voting in a real election
- A valid test requires knowing “ground truth”, that is, what the results of the election should be
- How do you know this in a real election?



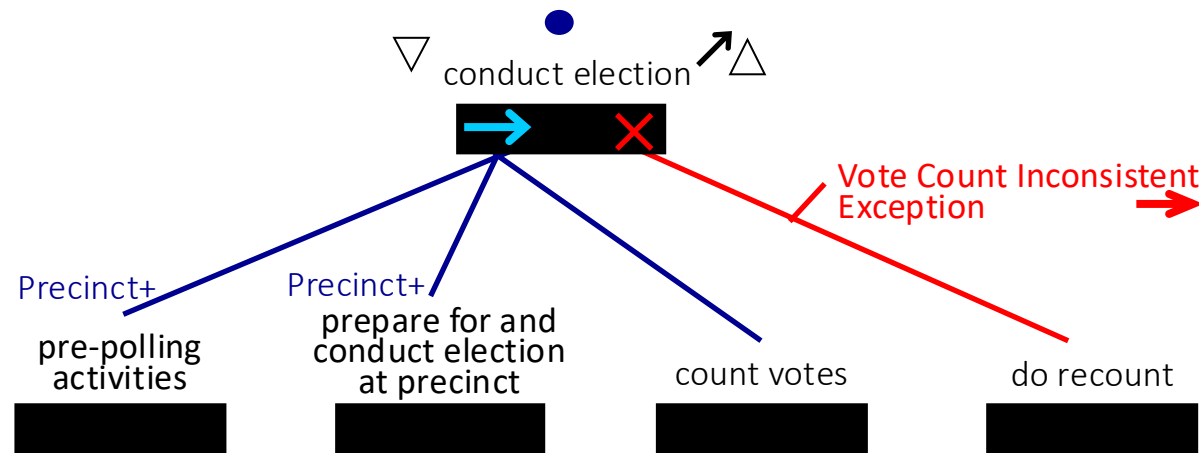
# The Take-Aways

- Know requirements of an election so you can *define* what you want
- Any computers used in an election process can be corrupted, so use good auditing techniques during the canvass
- And make sure the auditing techniques have good data!
  - Read: paper, as of now
- Given current election requirements, Internet voting poses great risks
  - The specific risks depend on how you do it

Remember, I don't have to rig an election to corrupt it; I just have to make you think I did!!!

# Election Process in Little-JIL

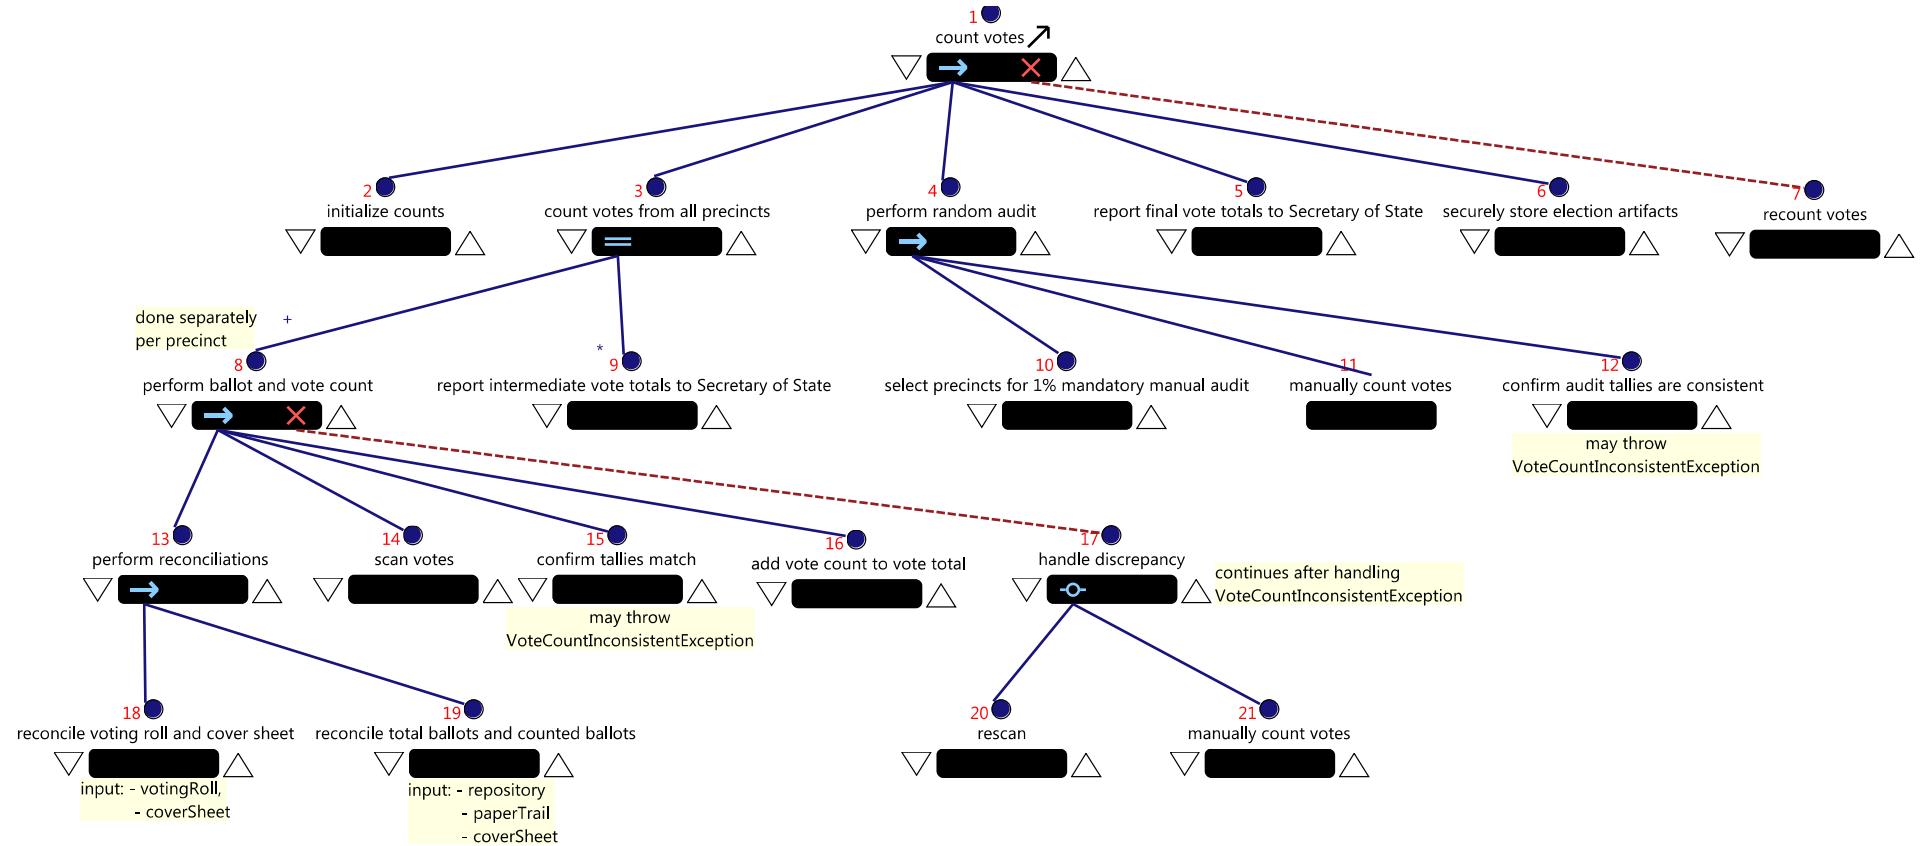
- Graphical process definition language with formal semantics; process represented as a hierarchical decomposition of steps



# Our Focus: Count Votes

1. Initialize counts
2. Count votes from all precincts
  - Count each precinct independently
3. Perform random audit
4. Report final vote totals to Secretary of State
5. Securely store election artifacts

# Subprocess “count votes”



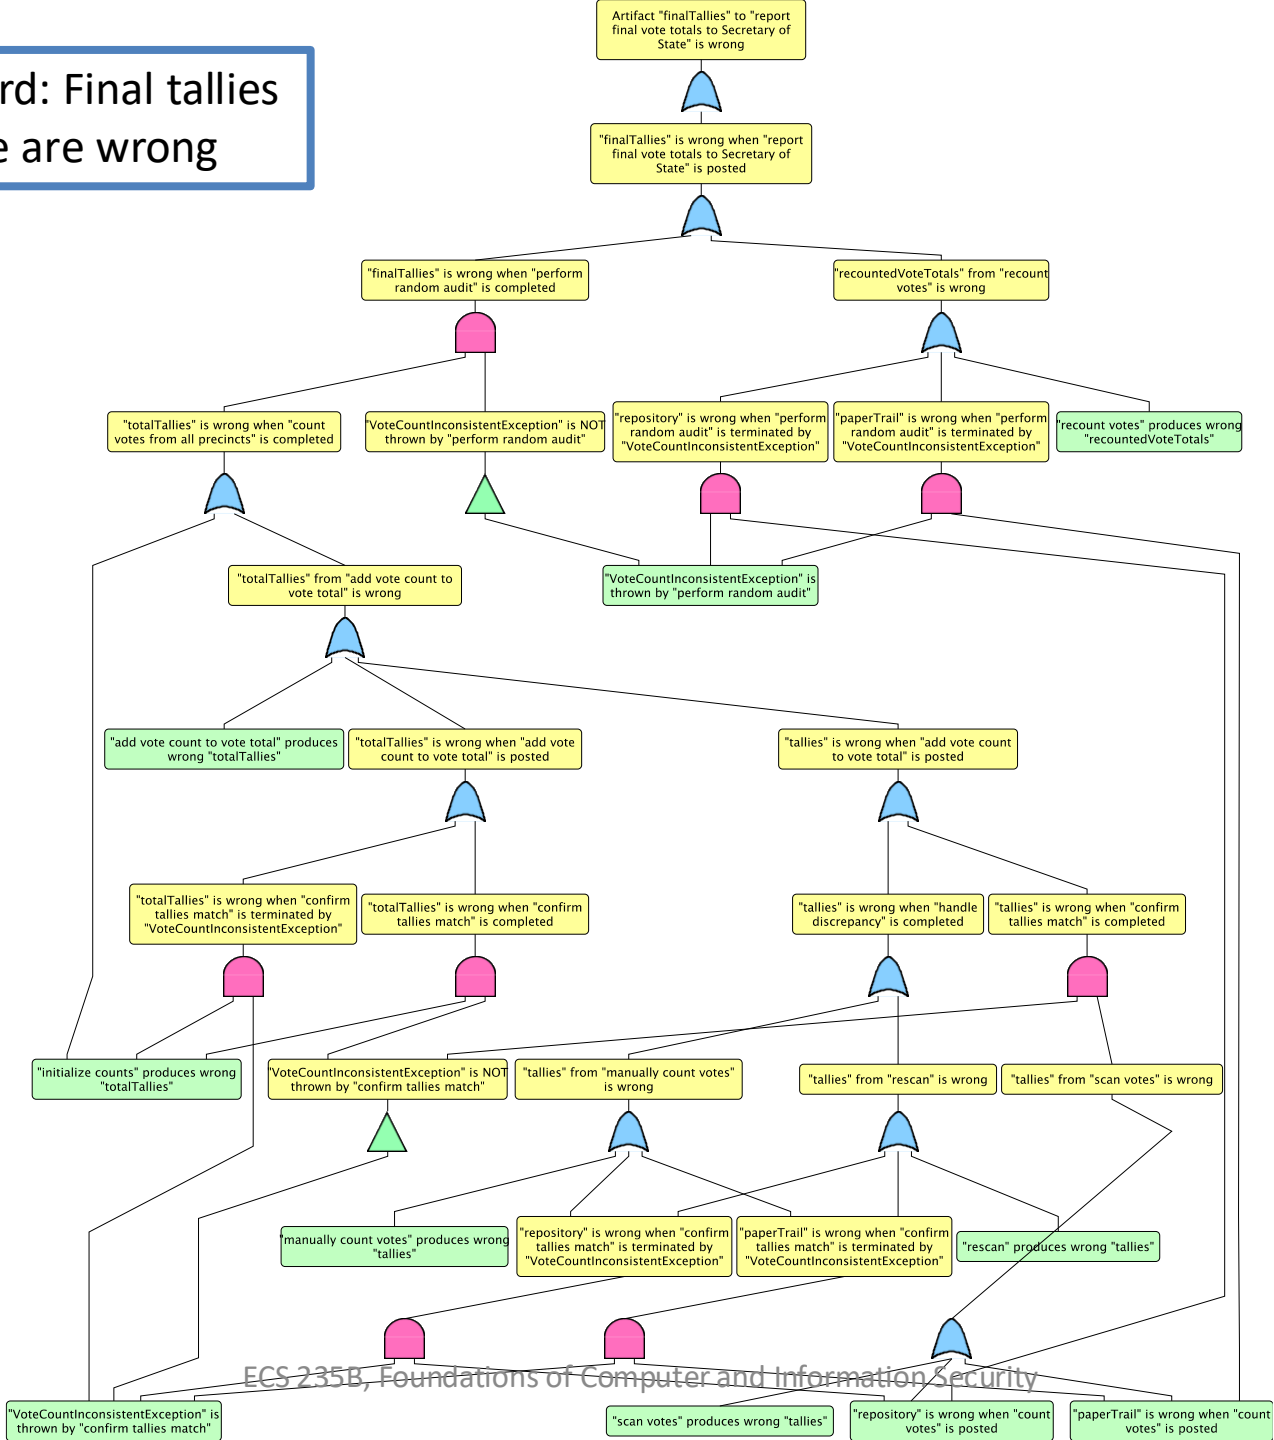
# Artifacts and Agents

(ref #) step	Input artifacts	output artifacts	agent
(2) Initialize counts		totalTallies	ElectionOfficial
(13) perform reconciliations	coverSheet; paperTrail; repository; votingRoll		ElectionOfficial
(18) reconcile voting roll and cover sheet	coverSheet; votingRoll		ElectionOfficial
(19) reconcile total ballots and counted ballots	coverSheet; paperTrail; repository		ElectionOfficial
(39) check off voter as voted	votingRoll	timeStamp	ElectionOfficial
(44) put ballot in repository	repository	timeStamp	ElectionOfficial

# Identifying Threats of Sabotage Attack

- Identify a *hazard* as the delivery of an incorrect artifact to a step in the process that delivers the artifact as a final process output
- From the process definition, automatically generate fault tree showing how hazard can occur

Selected hazard: Final tallies to Sec of State are wrong



# Example

- Hazard: wrong *finalTallies* delivered to the step *report final vote totals to Secretary of State*
  - Meaning the reported election results are wrong
- Automatically generate fault tree
- Use fault tree analysis tool to calculate minimal cut sets (MCSs)
  - Look for sets of activities where all agents are insiders and can modify final output (*finalTallies*) or artifact used to create final output

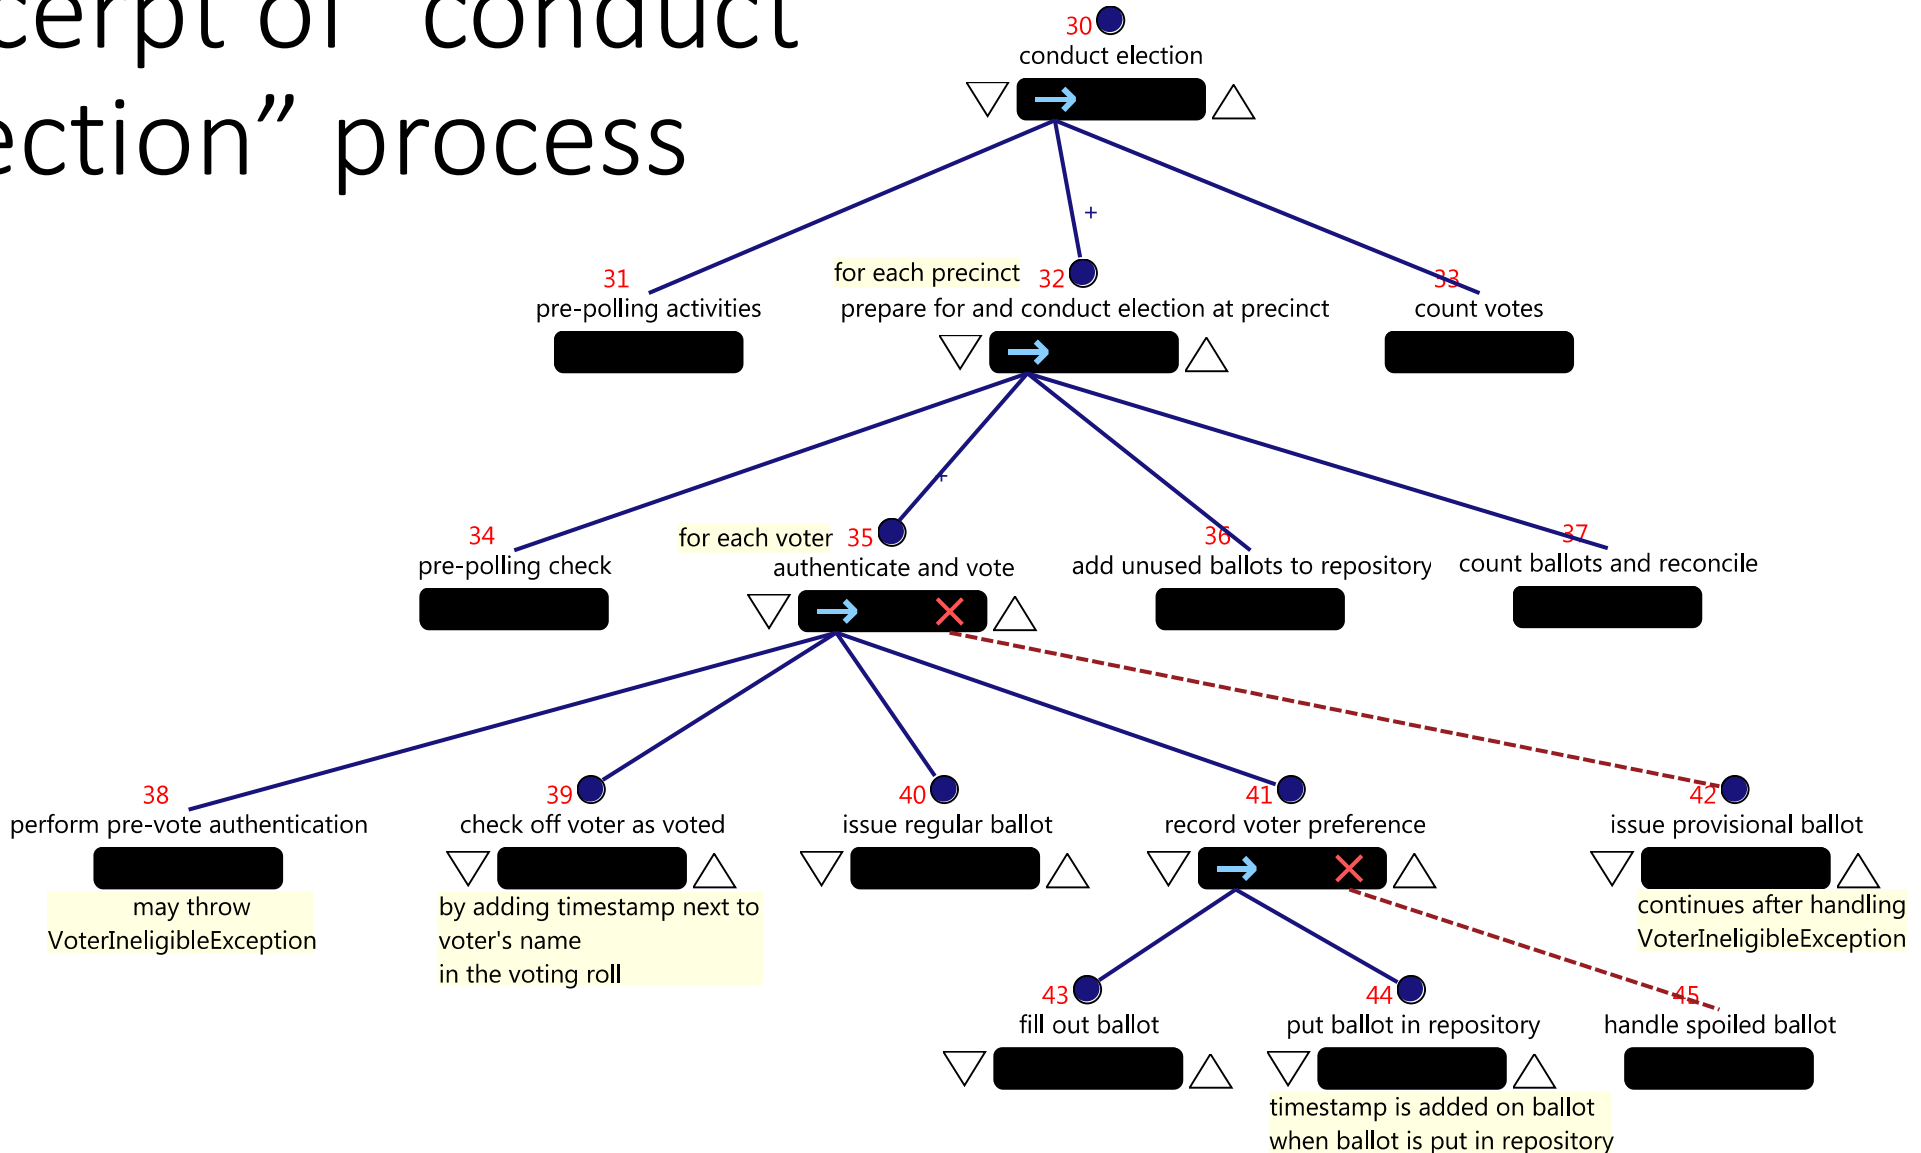
# 12 Possible Errors; Example Results

- 1**
  - Step *rescan* produces wrong artifact *tallies*
  - Step *perform random audit* does not throw exception *VoteCountInconsistentException*
  
- 2**
  - ▶ Step *scan votes* produces wrong artifact *tallies*
  - ▶ Step *confirm tallies match* does not throw exception *VoteCountInconsistentException*
  - ▶ Step *perform random audit* does not throw exception *VoteCountInconsistentException*
  
- 3**
  - ▶ Step *recount votes* produces wrong artifact *recountedVoteTotals*

# Data Exfiltration Attack

- In election context, associating a specific voter with a specific ballot
  - Done in Ohio, USA by correlating time-stamped ballots, poll books with times listed
- For expository purposes, voters vote on an electronic voting machine that time-stamps paper record of ballot
  - In Yolo, almost everyone uses paper, which is *never* time-stamped

# Excerpt of “conduct election” process



# Analysis

- If process executed as specified, only voter should know how she voted
- But . . .
  - Step 39: add timestamp next to name in roll
  - Step 44: add timestamp to ballot when placed in repository
- When can these be combined?
  - Artifacts are *votingRoll* (step 39), *repository* (step 44)
- Look in process model for a step, or sibling steps, using these artifacts
  - Steps 18, 19 here; parent is step 13, requiring both
  - *ElectionOfficial* is agent
  - So *ElectionOfficial* can exfiltrate data

# Evaluation

- Subjective
  - Process model validated by domain experts
  - Domain experts are better able to identify most worrisome types of insider attacks
- Objective
  - Focus on effectiveness, efficiency of process definition and analysis approaches
  - Little-JIL allows iterative process improvement based on feedback from domain experts

# Limitations

- Techniques are not always precise enough to fully describe the vulnerabilities and explain how they arise
- Analysis does not take into account full control and data dependencies of all steps
- Current agent descriptions are coarse
- Need to improve analysis of types of agents assigned to steps
- Use original analysis to suggest process modifications (automated or semi-automated)

# Conclusion

- Problem: instantiating the model
  - In particular, how do you get the ideal policy?
  - And how do you find the run-time policy?
- Need to determine threats
- How do you gather, analyze psychosocial information?
  - Social networks, media very useful for this
  - But one *heck* of an invasion of privacy!

# Key Points

- Treat attackers as a continuum, not as binary “inside” and “outside” divisions
- Policies aren’t precise, so think of them as layers of rules
  - Very useful for separating “intent” from “what’s actually implemented” at various levels
- Understand the *entire process*, not just the computer use!
  - Physical access often more important than computer access

# General Questions

- What operating system (and other third party software) does the e-voting system use, if any?
- How long does the e-voting system stay up?
- What is the procedure for getting e-voting machine ready to use?
- How have you tested the user interface? How do you handle write-in votes?

# General Questions

- How does the system associate votes with voters?
- Does your authentication/authorization mechanism associate external voter identities with that information?
- What is point at which e-ballot is cast, and voter cannot redo any part of the ballot?
- How do voters change their votes?
- How do you check enforcement of limits on voting in a race?

# General Questions

- What support must be provided to ensure that no-one can cast multiple ballots?
- What assumptions does the e-voting system make about procedures and support?
- How can the voter verify that the e-voting system accurately recorded votes cast, independent of the system (e.g., paper trails)?
- Does this verification require the intervention of a third party?

# General Questions

- What requirements is system designed to meet?
- How do you know it meets them?
- How do you handle updating software, hardware on fielded systems?
- How can you verify systems meets requirements after maintenance, upgrade?
- What audit mechanism, external vote tally, does the system supply?
- How could an auditor use this mechanism to validate results of an election?

# General Questions

- What happens if something goes wrong—do you lose (or change) votes?
- How do errors in setting up the system affect its ability to record and report votes?
- What happens if an election official or poll worker makes a mistake?
- How many mistakes (or malicious acts) will cause the results to be inaccurate?
- How many corrupt election officials involved with the systems are needed to corrupt the election?
- *Quis custodiet ipsos custodes?*

# General Questions

- Will the systems be available for testing?
- Will their components (software, source code, build environment, etc.) be available?
- Will any documentation required by law be available?
- Will the testers be able to reconstruct the system from what they are given, or will they have to take the vendor's word that the running system was in fact constructed from the given components?