RFID-Based Electronic Voting: What Could Possibly Go Wrong?

Yossef Oren and Avishai Wool
Computer and Network Security Lab
School of Electrical Engineering
Tel-Aviv University, Ramat Aviv 69978, Israel

***Abstract*—When Israel's Ministry of Internal Affairs decided to move to electronic voting, it chose to replace the traditional paper ballot with secure contactless smartcards. The system was designed around HF RFID technology to make voting stations easier to use and less prone to mechanical faults. However, in doing so the system was exposed to a powerful class of hardware-based attacks called relay attacks, which can extend the interrogation range of HF RFID tags far beyond the nominal range of 5 centimetres. We show how a low-budget adversary armed with a relay device can read out all votes already cast into the ballot box, suppress the votes of one or several voters, rewrite votes at will and even completely disqualify all votes in a single voting station. Our attacks are easy to mount, very difficult to detect, and compromise both the confidentiality and the integrity of the election system.**

I. Introduction

A. The Israeli e-Voting Scheme

This report discusses an e-voting system which is in the final stages of legal ratification in Israel[9] and is currently undergoing countrywide pilot testing. The scheme is officially described in a patent claim recently granted to the Government of Israel by the World International Property Organization [21], [4], as well as by the public tender to contractors implementing the scheme[8] and by the Israeli law governing the election process[24]. Further information was provided by discussions with Yoram Oren, one of the co-developers of the scheme[20] (not related to the authors).

The novelty of the system is that instead of using paper ballots, the votes in the proposed system are cast on contactless smartcards. To cast their votes, the voters use a computer terminal to write their choice into a contactless smartcard, and then physically deposit this smartcard into a ballot box. For a more detailed description of the scheme see section II. The designers of the Israeli e-voting scheme chose near-field contactless readers instead of traditional smartcards for non-security-related reasons. First and most important is the issue of cost and reliability – since a contactless smartcard reader has no mechanical interface and no moving parts (in contrast to a traditional smart card or magnetic-stripe reader), it can survive many more repeated uses with a reduced opportunity for damage or deliberate vandalism. In addition, as observed in [11], contactless smartcards are easier to use than magnetic stripe cards or traditional smart cards since they work regardless of the way the card is oriented with respect to the reader. Cost saving is also reportedly the reason why the system has absolutely no paper trail – the designers wished to save on the cost of maintaining and supplying paper to thousands of printers on election day.

The cards chosen for use in the Israeli scheme are Global Platform Java cards conforming to the ISO/IEC 14443[14] standard family. These tags contain no power source and rely on the reader to provide them with operating power. Due to the physical principles behind magnetic coupling, the reader is only capable of delivering this power to the tag if the tag is near the reader's antenna (Standard ISO/IEC 14443 readers and tags are designed for a 5cm operating range). This property, which imposes a severe limit on the usable range of the system, is sometimes considered erroneously as a security feature. To provide for the case of multiple tags sharing the same air space (such as communication with one out of several contactless smartcards which are all located in the same wallet), the ISO/IEC 14443 standard specifies an **anticollision** protocol which allows each tag to be interrogated in turn without corrupting the communications with other tags. An in depth introduction to RFID can be found in [7].

B. The Theory Behind Relay Attacks

The nominal 5cm operating range of ISO/IEC 14443 leads to the (mistaken) implicit assumption that whenever a reader can communicate with a tag one can assume that the tag is physically very close to the reader. As described in [16], relay attacks challenge this underlying assumption and allow a nearly unlimited distance between tag and reader. The attack achieves this ability by placing a **relay**, consisting of custom-designed tag and reader hardware connected by a high-range communication link, between the victim's tag and reader.

As illustrated in figure 1, the **leech** device (or proxy-reader) presents itself to the **victim tag** as a legitimate reader, while the **ghost** device (or proxy-token), presents itself to the **victim reader** as a legitimate tag. When the victim tag and reader wish to exchange data, their messages are captured by the relay devices and sent across

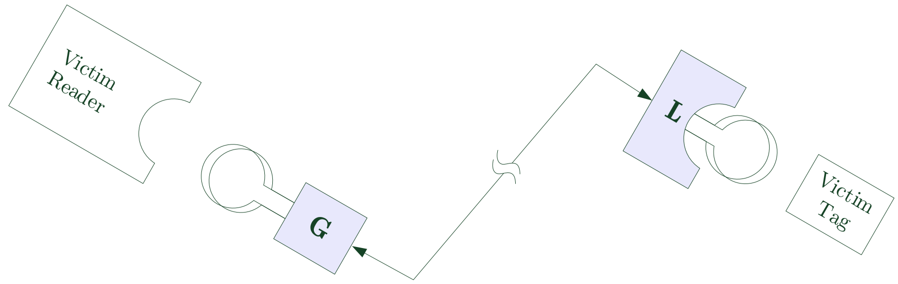

Figure 1. An RFID channel under a relay attack. Device “L” is the leech, while device “G” is the ghost.

the fast high-range communication link. Thus, any data exchange between the victim’s tag and reader can be carried out over the relay channel, even if it is strongly encrypted and authenticated. Such a relay attack on a credit card system, for instance, would allow the attacker to make a purchase and pay using an unsuspecting victim’s credit card. The attacker will present the ghost device to the point of sale system, while the leech would be placed near the victim’s wallet – perhaps in a totally different location.

Since the attacker’ ghost and leech hardware are custom manufactured and do not need to comply with any regulatory standard, the adversary has a large degree of freedom when designing his system. The adversary can specifically design his ghost and leech so that they have an internal power supply, a different form factor or (most importantly) a higher operating range.

Thus, there are three ways in which relay systems increase the distance between the victim tag and victim reader. They may increase the range between the **victim tag** and the **leech (leech range)**; they may also increase the range of the communications link between the **leech** and the **ghost (relay range)**; and finally, they may increase the range between the **ghost** and the **victim reader (ghost range)**

C. Previous work on Relay Attacks

An excellent survey on the history and state of the art on relay attacks has recently been written by Hancke et al.[11]. As stated in the survey, the first published discussion of a relay attack is arguably the “chess grandmaster” attack, which Conway writes about in his famous book “On Numbers and Games”[3].

In 2005 Kfir and Wool[16] first noted that applying a relay attack to a near-field contactless system will cause the security of such a system to “collapse”. They also used a detailed physical model of the physical layer to evaluate how the ghost and leech distances can be significantly improved beyond the nominal 10cm at a reasonable cost. They predicted that the leech range can be increased to 40 to 50cm and that the ghost range can be made as large as 50m. In that same year Hancke[12], [10] demonstrated a working relay attack on an NFC system using a high-speed radio link, using standard equipment for the ghost and leech radio interfaces and a low-cost radio transciever for the relay link. This setup provided a relay range of about 50 meters. In 2006 Kirschenbum and Wool[17] actually built a functioning leech element of the relay by demonstrating a low-cost RFID skimmer that provides a leech range of 25cm, thus validating the model-based prediction of [16]. The article also explained how the leech range can be improved even further while staying within a very low budget. A possible way to protect against a relay attack would be through the use of Faraday cage shielding, or by employing a hardware based distance-bounding protocol which strongly proves physical proximity (see [5]), but even these protocols cannot reliably detect a short-range relay attack.

In 2007 Drimer et al.[5] applied a relay attack to a contact-based smart card system implementing the high-security UK EMV payment system[6]. One original aspect of this system was the design of the attacker’s ghost and leech devices. Since the tag under attack was contact-based and designed to be used in a point-of-sale setting, the authors took special care to make the ghost look as much as possible like a standard EMV credit card. To make this attack undetectable in the field, the relay device and connecting cable could presumably be hidden up an attacker’s sleeve. The leech device (which is supposed to receive communications from the victim’s tag) was placed inside the box of a standard Chip & Pin terminal, resulting in a device that looks perfectly genuine to the victim.

II. A Detailed Description of the Israeli e-Voting Scheme

A. Components of the Scheme

The components of a voting station are illustrated in figure 2. Each voting station consists of the following elements:

- A **voting terminal** (a portable computer with a contactless smartcard reader). The voter uses this terminal to cast his vote. The vote is recorded twice: First, the individual vote is written onto the blank ballot (a contactless RFID smart card). Second, the total vote count is immediately tallied and this total is written to a regular contact-based smart card plugged into the voting terminal.

- A **verification terminal** (another portable computer with a contactless smartcard reader). This terminal is only capable of reading (and not writing) ballots. The voter can optionally place his written ballot on this terminal to make sure his vote was correctly cast.
- A set of **blank ballots**, taking the form of secure contactless smart cards which are cryptographically paired with this specific instance of voting and verification terminals (see subsection II-C).
- A **voting booth**, formed by folding a cardboard divider, that hides the voting and verification terminal from the elections committee and allows the voter to vote in privacy.
- A **ballot box,** where cast ballots (written contactless smartcards) are physically collected. The ballot boxes are typically made out of cardboard sheets, which are delivered flat and folded into shape on election day.
- The **local elections committee**, which typically consists of three mutually distrustful (and technologically inexperienced) members nominated by the parties participating in the elections. These local commitees in turn report to a central elections committee which oversees the election process. The committee uses a **population register terminal** to verify that each voter accessing the voting station is eligible to vote and has not voted before.

The election body has a very limited control over the physical environment in which the system is deployed – typically the voting station is located in a classroom or another location which is freely accessible to the adversary before election day. Furthermore, it is quite likely that one or more members of the local elections committee are in the service of one of the candidates (since the candidate parties nominate the members of these committees).

B. *The Proposed Voting Process*

The actual voting process is illustrated in figure 2, and is carried out in the following way:

1) The **voter** approaches the **elections committee**. The committee members verify his eligibility to vote, take his ID card (which is mandatory in Israel) and provide him with a **blank ballot** (contactless smartcard).
2) The voter enters the **voting booth**, where his actions cannot be seen by the committee. He next places his blank ballot on the reader connected to the **voting terminal**. He selects the vote he would like to cast via a touchscreen interface. Once the voter is satisfied with his vote it is written electronically to his ballot. The running total count of votes is also written to a (non-contactless) smart card embedded in the voting machine. The voter is allowed to change his mind and update his vote multiple times, with both the voting terminal's internal smart card and the contactless ballot tracking the latest choice.
3) If the voter wishes to convince himself that the vote he has selected was correctly recorded on his ballot, he may place the ballot on a reader connected to the **verification terminal**, which simply displays his selected vote. As noted in subsection II-C, the verification terminal is only capable of reading votes cast in this particular voting station. The ballot is supposed to be cryptographically secured and its contents cannot be read back by any other way.
4) The voter, now satisfied with his vote, drops his written contactless smartcard into the ballot box, in plain sight of the voting committee.
5) After witnessing the voting process, the elections committee returns the voter's ID card to the voter.
6) At the end of the day, the local elections committee retrieves the contact-based smart card from the voting terminal and delivers it to the central elections committee, where the cards collected countrywide are tallied to calculate the nationwide preliminary election results. These results are formed by adding together the running totals stored on the smart cards embedded in the voting terminals.
7) After the preliminary results are announced, the local elections committee manually counts all votes found inside the ballot box by passing them one by one through the verification terminal. As specified in [24], the final manual count takes precendence over the preliminary computer-counted results. If the "hand-counted" votes and the computer-counted votes disagree by more than 5 votes, the ballot is "flagged" and must be examined for tampering before being accepted in the total count. If the two tallies differ by more than 30%, all votes in this voting station are immediately ruled invalid. Any subsequent vote recounts, if desired, are performed against the smartcards inside the ballot box.

C. *Security Features of the Scheme*

The Israeli e-Voting Scheme was designed with a certain emphasis on security. The Global Platform Java cards used by the system conform to Common Criteria EAL 4+ [2] and are used in other high-security applications. The voting and verification terminals are cryptographically paired with the blank ballots used in each specific station, meaning that (at least as designed) a ballot cannot be read from, or written to, outside its specific voting terminal[1]. This means an attacker cannot steal a voting terminal from one voting station and use it to his advantage in another station. The voting terminals have no online connection either – the identity of the voter is only verified using by the population register terminal used by the voting committee. We have not reviewed these security features in any way.

[1]Even the government's "master key" is incapable of rewriting a ballot. It can only format the contactless smart card to a blank state

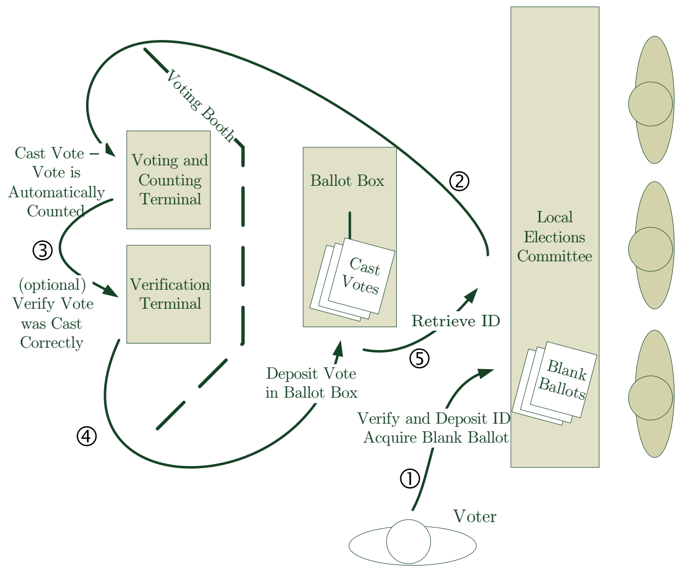

Figure 2. The proposed Israeli e-voting scheme in action. Illustrated from left to right are the voting booth, the cast ballot box and the local election commitee's desk area. The arrows show the path followed by a voter through the three areas of the voting station.

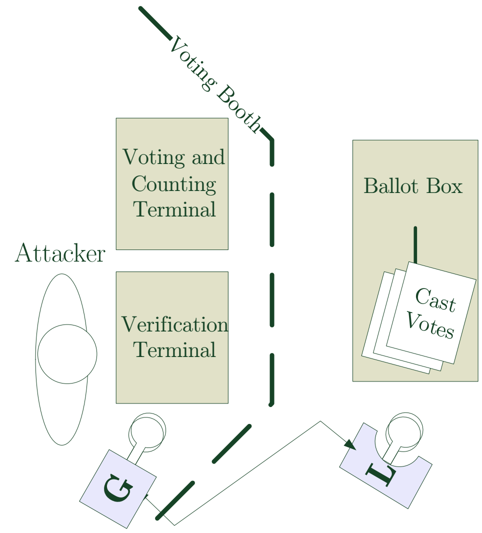

Figure 3. Ballot Sniffing Attack

The redundancy in the vote counting process offers another degree of security, since the voting tallies which are written to the secure smart card inside the voting terminal must match the count of votes in the ballot box. Thus, an attacker would theoretically need to subvert both locations before compromising the election results.

We must note that the specific protocol chosen to form the interface between tag and reader is of no interest to us. Indeed, as shown in [5], even the most advanced distance bounding protocols cannot defend against a properly implemented relay over a distance of less than 6 meters.

III. Relay Attacks on the e-Voting Scheme

As mentioned before, the fact that legitimate contactless smartcards have a limited read range should not be considered a security feature. Specifically, the fact that a certain tag is communicating with a certain reader does not imply with any certainty that the tag and reader are in proximity. In our attacks, we use this flaw to create a communications link from the voting and verification terminals inside the voting booth to the RFID-based ballots inside the ballot box that carry votes which were already cast.

While relay attacks have been discussed in other contexts, this is the first time they are considered within a voting system. One specific aspect of the voting process makes relay attacks even more effective in this context. In conventional relay attacks the attacker's tag (or ghost device) is generally used in a legitimate point-of-sale or conditional entry scenario, and as such must look very similar to an authentic tag to avoid suspicion. Fortunately for our attacker, this limitation on the physical form of the ghost does not exist when attacking the Israeli voting system – when performing relay attacks on the election terminals, the attacker is hidden inside the booth and is promised a high degree of privacy by the very nature of the voting process itself. This means that the ghost end of the attacker's relay setup can take an arbitrary form. The attacker may even use his time in the voting booth to replace the entire ensemble of voting and verification terminals with hacked machines running his own code (as the authors of [5] did when attacking EMV systems). The reader end (or leech device) can likewise be given an arbitrary shape, if we make the reasonable assumption that the voter can approach the voting commitee carrying a backpack or bag and ask them to mind his bag while he votes, or even consider the possibility that one of the three members of the elections committee is collaborating with the adversary.

The following subsections outline several relay attack scenarios that are possible against the proposed Israeli voting system.

A. The Ballot Sniffing Attack

This attack, as illustrated in figure 3, allows an adversary to learn the complete contents of all votes already cast into the ballot box.

1) Attack Description: To carry this attack, a relay is set up between the votes inside the ballot box and the verification terminal inside the voting booth. Next, the adversary repeatedly activates the verification terminal, each time with a different ballot. Since the verification terminal is paired with all cast votes in the ballot box, the adversary can now enumerate all votes inside the ballot box and display the votes recorded on them. This process can be carried out reliably and efficiently even if there are many votes already in the ballot box, as long as the "leech" component of the relay uses the RFID anticollision

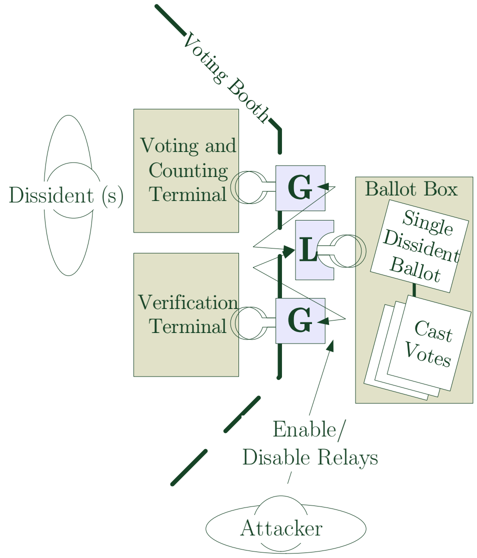

Figure 4. Single Dissident Attack

protocol to activate the tags one at a time[2].

2) Implications: Using this attack, an adversary can determine the partial results of the vote before the voting day is over. If the votes in a certain ballot do not fit the outcome expected by the adversary, he may use this data to decide to disqualify the entire voting terminal using other methods described below. More importantly, if this attack is carried out twice in one day on the same ballot box it will teach the adversary about the votes cast in the interval between the two attacks. Thus, this attack can be used to verify that a voter had voted a certain way, violating voter privacy and allowing coercion to be inroduced into the voting process.

B. The Single Dissident Attack

This attack, as illustrated in figure 4, lets the adversary suppress nearly all undesirable votes in a certain voting station while remaining virtually undetectable.

1) Attack Description: To carry out this attack, the adversary needs to be present near the voting terminal while another voter is voting (for example, the adversary may volunteer to serve on this specific election committee or simply stand outside). As a prerequisite to this attack, the adversary sets up a relay between the voting and verification terminals and a single previously cast vote located inside the ballot box, hereby called the "single dissident vote". This relay can be turned on and off as desired by the adversary. The adversary then selects a certain group of voters whose voice he wishes to suppress[3]. In general, the relays will operate only when one of these "dissidents" is inside the voting booth. As stated previously, an attacker can easily obtain such fine-grained control over the voting and verification systems if he replaces both terminals inside the voting booth with similar-looking computers running his own code (see [5]).

When the relay attack is active, all of the dissenting voter's actions (voting, verifying, etc.) are not performed against the blank ballot he is holding, but rather shunted to the single dissident vote already in the ballot box. Since this single ballot is properly authenticated, both the voting and the verification terminals happily register the vote and display its correct value. Thus, the voter has no idea that the attack is taking place, but his personal ballot always stays blank. After the voter exits the booth, he casts his blank vote into the ballot box and his vote is effectively disqualified and ignored.

2) Implications: If the adversary manages to correctly guess the disposition of the voter, this attack makes sure the ballot box contains no more than one dissenting vote for each undesirable party. Because of the relay attack all other dissenting voters will register as blank, or invalid, votes, and have no effect on the outcome of the elections. Since the votes inside the ballot box correlate directly with the values written to the smart card inside the voting terminal, the recount at the end of the voting day will find no discrepancy between the two. However, there is only one vote left to represent any amount of dissenting voters.

While this attack seems at first difficult to set up and carry out, it has the advantage of being virtually undetectable. If one of the dissenting voters gets suspicious and tries to find traces of this attack, even manual examination of the per-box vote counts will not reveal the subterfuge in this case, since each voter can be certain only of his personal vote, and at least one dissenting vote has been registered in this ballot. Thus, each dissident will be forced to conclude that his accomplices had a change of heart at the ballot, and that he is indeed the only one to have placed a dissenting vote in this specific ballot box.

If the attacker is somewhat uncertain of his ability to guess how a voter will behave inside the booth, he can use not one but several dissident votes and map the dissenting voters at random into one of this small set whenever one of them enters the booth. Since only a single vote for each party needs to be registered at the voting terminal for the attack to be undetectable, the probability of failure can be significantly reduced by choosing an appropriate size for this group.

The implementers of the system may try to keep track

[2]The ISO 14443 anticollision protocol is designed to discriminate between up to 2^{32} tags. However, according to [20] it has some problems in practice when dealing with more than 3 tags at the same time. If indeed this is the case, this may limit the effectiveness of the ballot sniffing attack.

[3]The attack builds on the assumption that the the adversary can guess the vote of a some voters based on their external characteristics or some other auxiliary information. This is very reasonable in general, and even more so in Israel given the highly heterogenous nature of the Israeli voting body.

of which ballots were written to, and refuse to connect to previously written tags as soon as a new tag is encountered[20]. However, it will be difficult to provide this functionality without exposing the system to voter privacy violation or to denial of service attacks.

C. The Ballot Stuffing Attack

This attack gives an adversary complete control over previously cast votes, using a relay attack to rewrite them to the candidate of his choice.

1) Attack Description: In this attack the relay is set up between the ballot box and the voting terminal. The adversary then enters the voting booth and repeatedly votes for the candidate of his choice, each time engaging with a different ballot selected from the ballots already inside the ballot box. Since the voting terminal is paired with the previously cast votes, it has no reason to prevent the votes from being modified. Since the voting terminal itself is used to perform the attack, it constantly updates its running vote counts and thus remains in perfect sync with the cast votes instead the ballot box[4]. This means that this attack causes no discrepancy which can be detected during a recount. The RFID anticollision protocol allows this process to be carried out reliably and efficiently even if there are many votes already in the ballot box. The ability to rewrite ballots multiple times is a **confirmed feature** of the system which is stated in the voting law[24].

2) Implications: This attack is the most conspicouous of all attacks presented here, but it is the most powerful, since it allows the entire set of ballots to be rewritten arbitrarily. As noted in the previous subsection, the only sort of fraud that is detectable by mutually distrustful voters would be the case when a party the voter had personally voted for ends up with zero votes in this ballot. If the adversary had previously applied the ballot sniffing attack to read all votes in the ballot, he can avoid this form of detection by registering one vote for each party originally represented in the ballot under attack and giving all other votes to his candidate of choice. The lack of paper trail means that this attack is undetectable and unprovable, even though it has the potential to arouse voter suspicion if used crudely. Note that the countermeasures sketched in the two previous subsections affect this attack as well.

IV. Non-Relay Attacks

The relay attacks described in the previous section require some sophistication and a minor budget from the attacker. In addition to these attacks, the RFID technology used in the proposed voting system is also susceptible to simpler hardware-based attacks.

A. The Zapper Attack

This relatively low-tech attack can quickly and easily disqualify a certain ballot box, allowing an adversary coarse-grained control over the results of the election.

1) Attack Description: The adversary can make use of a low-cost "RFID Zapper" device[18], consisting of a pulsed power source (most conveniently, the flash circuit from a disposable camera) connected to a properly shaped reader antenna. When the pulsed power source (e.g. flash camera) is activated, a powerful electromagnetic pulse flows through the antenna and "zaps" any RFID tag close enough to couple magnetically with this antenna with a powerful surge of current which overwhelms its input circuits and generally renders it unusable. A live video demonstration of such a device being used to disable an ISO/IEC 14443 tag can be found in [1, starting at 19:00]. Once enough tags inside the ballot box are "zapped"[5], this will cause a discrepancy between the count of ballots held by the smartcard located inside the voting terminal and those counted manually by the elections committee, causing the entire ballot to be legally disqualified.

Zapping attacks against a variety of ISO/IEC 14443-compatible HF tags were experimenatally verified by several research groups [23], [19]. The zapped tags included low-cost inventory tags as well as high-security tags similar to the ones used in the Israeli elections system[6]

2) Implications: This attack allows anybody with access to a ballot box to quickly and undetectably disqualify its contents and remove it from the election process. The ballot box can be zapped during the election process or later, while the ballots are kept pending the manual recount. In a particularly clever twist on this attack, a small group of crafty adversaries can zap their *own* cards before depositing them in the ballot box, thus immediately disqualifying the entire ballot (of course, the privacy features inherent to the voting process will prevent these individuals from being tracked or caught). The zapping attack is so easy and inexpensive to carry out that it may be even done for reasons of malice or protest, and not just for changing the election results.

B. Jamming, Blocking and Selective Denial of Service

Jamming attacks can disrupt the operation of the RFID reader at a distance. They can be turned on and off on demand.

1) Attack Description: While the physical interface used by the tag is based on magnetic coupling, most RFID readers actually use a standard radio decoding technique called single side-band (SSB) demodulation to decode the tag responses. This property allows the ghost component of a relay to work at distances of up to 50 meters, but it also means that the reader's standard operation can be disrupted by SSB modulated signals sent to it from

[4] We assume that the voting terminal handles a voter's vote changes by rewriting the ballot, subtracting one vote from the old choice and adding one vote to the new choice on the internal smartcard, and that there is no "finalizing" step which locks the ballot from subsequent edits. This interpretation was validated by [20].

[5] At least five tags, and at any case no more than 30%[24]

[6] Specifically, an admission ticket to the 2006 FIFA world cup.

a distant radio antenna. To carry out this attack the attacker uses a directed antenna to transmit a random SSB-modulated bit pattern toward the reader at the tag's upper side-band frequency of 14.4075 MHz. The reader connected to the voting terminal will not be able to distinguish between this disruptive signal and the tag's response, rendering the reader inoperative and creating an effective denial of service attack. According to [16] the operational range of such an attack can be up to 50 meters.

A similar attack can be performed by hiding a "Blocker Tag"[15] somewhere within the voting booth (for example, between the voting terminal and the table on which it is set). This tag is a special hardware device which conforms only partially to the RFID specification – specifically, it ignores the tag anti-collision protocol under certain conditions, disrupting tag and reader communications. The blocker tag presented in [15] prevents RFID readers from communicating with tags whose ID matches certain criteria while allowing all other communications. The design can be modified to support a high-range remote control interface allowing blocking to be turned on and off on demand, creating another hard to detect denial of service system.

2) Implications: The jamming attack is a **selective denial of service attack**, since it is easy to apply selectively only to a certain subset of voters at the discretion of the attacker. This attack is also unique in its very high operating range. It cannot be prevented unless electromagnetic shielding is applied to the walls, doors and windows of every voting station (and not just the ballot box itself).

C. Fault Attacks

As demonstrated in [13], HF RFID tags are highly susceptible to Optical Beam Induced Current (OBIC) attacks, which belong to the general category of of *fault attack*. When subjected to this attack, the tags' writing mechanism becomes completely unreliable. In the case of HF RFID tags, this attack is not at all invasive or expensive, requiring no manipulation of the tag other than shining a high-power LED at the its chip area. Such a LED can be easily hid by the adversary inside the voting terminal. Quoting from [13], the authors show that "RFID tags are exceedingly vulnerable to faults during the writing of data that is stored into the internal memory [...] it is possible to prevent the writing of this data as well as to allow the writing of faulty values. In both cases, tags confirm the operation to be successful."

Such an attack can have a devastating effect on the described voting system, since a failed write which is reported to the voting terminal as successful immediately translates to a discrepancy between the voting station's tally and the final manual count, causing in turn the entire voting station to be disqualified during the manual recount phase.

V. Countermeasures

This section lists several countermeasures which can be applied to increase the security of the proposed system. Even with all of these countermeasures applied, the system as proposed seems "too broken" to be used in a high-stakes democratic process. However, it would make good sense to apply them if a future variant of the scheme is used for a less important purpose than a democratic election.

A. Physical Layer Security

The immediate and most crucial protection mechanism against relay attacks is to shield the voting station and its environment from all forms of electromagnetic radiation. Most importantly, the ballot box should be made into a Faraday cage by constructing it from a sufficiently thick sheet of conducting material such as aluminum. Ferric materials such as iron are probably less suitable since they may not block magnetic coupling between reader and tag. Unfortunately, it was shown that RF shielding against relay attacks is difficult in practice to do properly and cheaply [22], especially for ballot boxes which should be portable, foldable and lightweight. RF shielding is also made difficult by the fact that the ballot box should have a slot for accepting votes[7]. More troubling is the fact that electromagnetic shielding should also be applied to the walls, windows and doors of the voting station itself, to prevent jamming, zapping and selective denial of service attacks from being carried out from a distance. A handy indicator of proper shielding would be the total lack of cellphone reception inside the voting station. When outside the ballot box, the individual ballots (used or not) should be carried inside envelopes made of conducting metal.

The ballot box should also be kept away from any object large enough to hide a leech device. To be on the safe side, no object should be allowed to come within a 1 meter radius of the ballot box. In particular, the ballot box should **not** be on or near the committee table. One open question is how to allow an untrusted voter to approach the ballot box and deposit his vote.

B. Audio Feedback and Cooldown

The voting and verification terminals should be programmed to emit a loud beep each time they read or write a tag. In addition, the terminals should have a "cooldown" period of at least 30 seconds between each use, or at least before switching from one tag to another. Any attempts at wholesale vote sniffing or rewriting will now take a long time and result in many conspicuous beeps, a fact that should arouse the suspicion of the voting commitee. Note that these countermeasures have no effect on the single dissident attack (which registers a single vote per voter),

[7] A possible approach would be to protect the slot by a catflap-type arrangement or to line it with fine metal bristles which touch when no card is present.

nor do they prevent a single voter from modifying a small but constant amount of votes.

An additional form of audio feedback can be implemented by prepopulating the ballot box with a "canary" tag whose sole functionality would be to beep when interrogated. Since votes inside the ballot box are not typically expected to be interrogated, any attempts to perform a relay attack without first dealing with the "canary" tag will be quickly detected.

C. Single-Write Ballots

The voting station should include some sort of irreversible "commitment" or ballot cancellation process (similar to the one performed on paper stamps) which will finalize the vote so the RFID smartcard can be read from but not written to again. In the best case this can be accompanied by a visible mark on the ballot itself. The single dissident attack described in subsection III-B will still work in this case if the adversary belongs to the elections committee, since he can keep one un-committed smart card especially for this purpose and perform the relay attacks against this dedicated card.

VI. Conclusion

The designers of the voting system most probably chose contactless technology for cost and reliability reasons. However, this choice led to a devastating consequence in the form of the relay attack. In this work we have shown how any party with moderate technical expertise and a fairly small budget can completely compromise the new Israeli e-voting system. Our attacks are aided by the fact that the voter's actions in the voting booth are unmonitored, by the design choice of making the tags rewritable, and by the additional fact that the local elections commitee are typically technologically inexperienced.

Despite the threat of relay attacks, contactless RFID smardcards are in use in other sensitive applications such as credit cards, e-passports and secure entrance systems, so one may think that they are good enough for e-voting as well. Unfortunately, in these other applications there are always additional security mechanisms: audit trails, credit card statements, insurance, security cameras, human guards, store clerks or border officials. In a voting application, voter privacy **by design** eliminates the possibility for all of these additional safeguards. Unless the attacker is caught "red handed" during the attack, a relay attack on an election may well be a perfect crime.

In its current form, the proposed system cannot guarantee the privacy of voters, and it cannot promise that cast votes were not manipulated after the fact. Even if countermeasures are properly implemented, we believe the proposed system is strictly inferior to, and completely unacceptable as a replacement for, the plain-paper ballot system currently used in Israel.

References

[1] CCC-TV lightning talks day 1. Online, 2005.

[2] Common Criteria Recognition Agreement. Common criteria for information technology security evaluation part 2: Security functional components. Online, July 2009.

[3] J. H. Conway. *On Numbers and Games*. Academic Press, 1976.

[4] B. Dolev. Laying the groundwork for electronic elections in Israel (in Hebrew). Invited Talk, CPIIS IDC/TAU Workshop on Electronic Voting, May 2009.

[5] S. Drimer and S. J. Murdoch. Keep your enemies close: distance bounding against smartcard relay attacks. In *Proceedings of 16th USENIX Security Symposium*, pages 1–16, Boston, MA, 2007. USENIX Association.

[6] EMVCo contactless communication protocol specification v2.0.1. Online, July 2009.

[7] K. Finkenzeller. *RFID Handbook : Fundamentals and Applications in Contactless Smart Cards and Identification*. John Wiley & Sons, 2003.

[8] Government of Israel, Ministry of the Interior. Public tender 16-2008 for the establishment and operation of a computerized election system. Online, August 2008.

[9] Government of Israel, Prime Minister's Office. Decisions of the ministerial committee on legislation. Online, August 2009.

[10] G. Hancke. A Practical Relay Attack on ISO 14443 Proximity Cards. Manuscript, February 2005.

[11] G. Hancke, K. Mayes, and K. Markantonakis. Confidence in smart token proximity: Relay attacks revisited. *Computers & Security*, In Press, Accepted Manuscript:–, 2009.

[12] G. P. Hancke. Practical attacks on proximity identification systems (short paper). In *SP '06: Proceedings of the 2006 IEEE Symposium on Security and Privacy*, pages 328–333, Oakland, CA, 2006. IEEE Computer Society.

[13] M. Hutter, J.-M. Schmidt, and T. Plos. Rfid and its vulnerability to faults. In *CHES '08: Proceeding sof the 10th international workshop on Cryptographic Hardware and Embedded Systems*, pages 363–379, Berlin, Heidelberg, 2008. Springer-Verlag.

[14] International Organization for Standardization, Geneva. *ISO/IEC 14443-2 Identification cards – Contactless integrated circuit(s) cards – Proximity cards – Part 2: Radio frequency power and signal interface*, 2001.

[15] A. Juels, R. L. Rivest, and M. Szydlo. The blocker tag: selective blocking of rfid tags for consumer privacy. In *CCS '03: Proceedings of the 10th ACM conference on Computer and communications security*, pages 103–111. ACM Press, 2003.

[16] Z. Kfir and A. Wool. Picking virtual pockets using relay attacks on contactless smartcards. In *International Conference on Security and Privacy for Emerging Areas in Communications Networks*, pages 47–58, Los Alamitos, CA, USA, 2005. IEEE Computer Society.

[17] I. Kirschenbaum and A. Wool. How to build a low-cost, extended-range RFID skimmer. In *Proceedings of the 15th USENIX Security Symposium*, Vancouver, B.C., Canada, 2006. USENIX Association.

[18] T. "Minime" and C. "Mahajivana". RFID zapper. 22nd Chaos Communication Congress, December 2005.

[19] N3ldan. Cheap/free capacitor bank and charger. Online, October 2006.

[20] Y. A. Oren. Personal communication.

[21] Y. A. Oren, P. Rosenblum, O. Margoninsky, I. Yom-Tov, and B. Dolev. (wo/2010/010564) electronic voting system. Online, January 2010.

[22] T. Plos, M. Hutter, and C. Herbst. Enhancing side-channel analysis with low-cost shielding techniques. In C. Lackner, T. Ostermann, M. Sams, and R. Spilka, editors, *Proceedings of Austrochip 2008, Linz, Austria, October 8, 2008, Proceedings*, pages 90 – 95, 2008.

[23] T. Runge. Schriftliche arbeit jugend forscht: Der rfid-zapper. Online, February 2007.

[24] M. Shitrit. Local authorities bill (elections) (amendment - election systems). Online, May 2009.