Putting the Human Back in Voting Protocols
(Transcript of Discussion)

Peter Y.A. Ryan

School of Computing Science, Newcastle University

I'd like to talk about the role of the human in voting protocols. Basically I want to argue that voting protocols seem to be particularly interesting from the point of view of the theme of this workshop, in the sense that the users of the system actually play a particularly important role in trying to maintain the assurance of the system itself. I'm interested in a particular class of voting protocols, so-called voter verifiable schemes, which aim to allow the voter to play an active role in contributing to the dependability and assurance of the system. In designing these systems, clearly that we want high assurance of accuracy, but on the other hand we have to balance that with maintaining the ballot secrecy, so that nobody can work out which way a particular individual voter voted, and we want to do it in such a way that we place minimal, or ideally, zero trust in components, such as, hardware, software, the voting officials, and so on, and suppliers.

There are some down sides to the approach that we've proposed. This kind of authority that we've postulated has to be trusted; not for accuracy, because the auditing means that we don't have to trust them for accuracy, but we still have to trust them for secrecy, to keep this material that they've created secret. We need to place some trust in the auditors; now arguably, of course, the trust is minimal, particularly if we've got a set of conflicting hostile auditors, we can spread the trust. We also have a chain of custody issues, we have to keep this material secret, because otherwise we get into chain voting coercion type attacks if some third party can get hold of the ballot form material ahead of time.

Bruce Christianson: Can you describe what you mean by a chain?

Reply: The idea actually works with conventional paper ballot forms, but it's particularly virulent here. The idea is that some adversary coercer gets hold of a blank form, the coercer marks it with a choice of candidate they want and as voters are coming into the polling station the coercer presents the voter with one of these forms and says, if you come out with a blank form I'll give you ten bucks, or whatever it is. Once you've started this, as long as the voters go along, you can continue indefinitely into the future. In the UK system, the voter registers, gets a fresh form, and then they supposedly go off to the booth and mark it, and cast it. And in this attack, the voter is being induced to cast the ballot form which was marked by the coercer, and come out with a fresh ballot form that they were given when they registered, and this kind of attack is particularly virulent here because of the web bulletin board. So here the danger would be if a coercer managed to get hold of one of these forms, and so they know the

B. Christianson et al. (Eds.): Security Protocols 2006, LNCS 5087, pp. 20–25, 2009.

association of candidate list and cryptographic number variable on the ballot form, they can of course later visit the web bulletin board and check that the voter really did use that ballot form and mark the cross in the correct position. Does that make sense?

Michael Roe: So for example if the coercer goes to the polling station first, has some piece of paper about the right shape with them, and by sleight of hand puts it in the ballot box, that means he's still got in his pocket the real ballot form, to start the chain.

Reply: Yes, so that's how you would initialise that kind of attack, or you bribe an official or something. There's always some way you can initialise the attack.

Bruno Crispo: This kind of attack is very peculiar to the UK, because in another place you cannot do that.

Reply: It turns out in France and Greece there's a different system. The problem with the UK system is that the ballot forms are a controlled resource, and so if you come out with a blank form, that suggests to the coercer that you did cast the marked form. Now as I understand it, in France there's a difference; the ballot forms basically lie around, and you actually register at the point that you cast your vote. Is that right George?

George Danezis: Yes.

Reply: Which actually helps, certainly with the conventional pen and paper systems. But actually if you think about it, it doesn't really help you with these cryptographic schemes, because the attack still works since the ballot forms have unique numbers, etc, etc. So there is a countermeasure which works for ordinary pen and paper systems, but it fails for these cryptographic systems. Perhaps we should distribute the creation of the ballot forms in such a way that no single entity knows the secrets, and in such a way that we just reveal, say, the candidate list, at the last moment in the booth so as to avoid these chain of custody, chain voting, attacks. In a sense, we want on-demand creation of the ballot forms, in the booth, so only the voter sees the candidate list at the time that they need it, and the candidate list then should get destroyed immediately they've made their mark on the form.

The difficulty now is: how do we ensure that the device in the booth shows the voter the correct candidate list? Previously we solved that problem by having the pre-commitments and the pre-auditing of the ballot forms, if we're doing an on-demand creation, of course, we can't exploit that anymore.

So that seems to suggest, as far as I can see, that we have to go back to a kind of "cut and choose", or similar tactic, to try and detect any corruption by the device in the booth. There seem to be several possibilities, but I'll just describe one here, which is actually to print double-sided ballot forms in the booth. So in effect these are two independent Prêt à Voter forms that you saw earlier, but printed on flip sides of a sheet. The voter can randomly choose one side to use to cast their vote, so let's suppose in this case the voter's chosen the left-hand side, they've put a cross against Plato. They should leave the other side untouched,

and in accordance with the previous system we again have to destroy the chosen side to cast. So we end up with a receipt which has the two sides; the left-hand side is the receipt of the form that you saw earlier, but on the flip side we set things up in such a way that the ballot form remains intact with the candidate list, and the crypto material here, and the point of doing that is that this can then be audited and checked at a later stage, or later stages, in fact. So the device, if it were trying to cheat, would in some sense have to predict which side the voter was going to choose, and perhaps corrupt that side, and then make sure that the other audited side was correctly constructed.

Bruce Christianson: Does this give me a 50% chance if I'm trying to corrupt the system?

Reply: Of buying the vote? I don't think it helps with buying. The point about this "cut and choose" thing depends what kind of attack you're worried about. There's privacy and accuracy, and this is primarily against accuracy, trying to guard against the device being able, in an undetected way, to corrupt the construction of the ballot forms in such a way that the vote would be incorrectly decrypted at the end.

Bruce Christianson: Ah right, I understand, if it's doing it systematically...

Reply: Yes, that's right, in each case it's got a 50-50 chance of getting away undetected, but of course if it tries multiple times, it falls off exponentially.

Matt Blaze: So the receipt you walk out of the booth with, and the cryptographic thing on the other side, that number of the second ballot, that doesn't help them construct a new ballot to vote again with. The Plato ballot is invalid?

Reply: This should get invalidated, yes. There are issues about how you make sure that ballots can't be recast, and how, for example, you make sure that you can only audit the correct side, because of course you don't want to start revealing seed information. So there are issues there.

Matt Blaze: So, you can't use this receipt, in an obvious way, to learn who someone voted for, but you can learn whether, for example, they submitted a valid ballot. Can I pay people for spoilt ballots, with two Xs, if I want to pay people to *not* vote?

Reply: That's an interesting point, and yes, people do worry about that; this is a sort of coercion to abstain. An issue is whether you have some sort of device in the booth to kind of assist the voter in this process, because this is now getting to be a more subtle process than I described previously. So you could, in theory, have a device which allows you to put just one cross in one side and which automatically destroys the left-hand side. That might help address those kinds of problem, but of course it throws up a whole host of other problems: do you trust this thing which assists the voter, and so forth. In a sense the message I'm trying to get across is there are trade-offs that we have to play against here. And

in the remote context things shift again, the issues are completely different, and there's a different sort of answer to your problem.

James Heather: But you can coerce somebody into not voting anyway because you can coerce somebody into not turning up to the ballot station. You don't have to force somebody to cast a spoiled paper.

Matt Blaze: But that's harder because I have to watch you all day as opposed to simply paying you.

Bruce Christianson: You can imagine adding a candidate called "none of the above" to the ballot paper. Then you can tell the difference between someone who's exercised their democratic right not to cast a vote, and someone who has spoiled the paper.

Chris Mitchell: It doesn't help to persuade people not to turn up if they're legally obliged to vote, because in order to coerce them into not voting, you do need a mechanism to be able to discover whether they've deliberately abstained.

James Heather: In places where you're legally obliged to vote, are you are still allowed to spoil the paper?

Bruce Christianson: Yes, although in Australia, where every citizen is legally obliged to vote, candidates whose names begin with A are greatly sought by political parties.

Reply: Incidentally, that's a sort spin-off advantage of this kind of scheme because everything's randomised each time.

It seems to me voting systems are peculiar in the sense that we really would like to set things up so that the assurance arises ultimately from the users themselves. Whereas we tend to think of the users as being the weak link, here we're trying to build the whole structure of trust on this bed of sand, the weak links, the voters themselves. That makes things very tricky, particularly with the electorate at large: just how much can we trust their understanding and diligence in executing the protocols? So it really is a very challenging sort of design problem.

Bruce Christianson: If you can get this politician elected we will pay you two hundred guineas.

James Malcolm: We're all very used to systems becoming more and more complicated, and more and more difficult to understand. Are there, following on from Matt's talk[1], any ways you can see in which to measure and evaluate the trade-offs and decide whether it worth putting in this extra complexity?

Reply: Well, that's a good question. I suspect you probably will have to run trials to try and evaluate voter reaction to some of these trade-offs. And there's a whole host of other issues surrounding the extent to which will people understand and trust the systems, and be prepared to use them. And there are issues about metrics as well, as prompted by Matt's comment, which I think are going

[1] Blaze, these proceedings.

to be particularly challenging with voting systems because it's a whole socio-technical system here. Quite how you evaluate the possibility of certain patterns of collusion, for example, in a metric, is going to be really tricky, so there's a whole bunch of interesting challenges here. I don't quite know how you evaluate these things, it's a major issue and I'd welcome ideas.

Mike Bond: On the idea of this trade-off between what benefit you get vis-a-vis the extra complexity of that, my opinion is that building coercion resistance ultimately is going to be entering into an arms race which the voting scheme manufacturers are always going to lose. The reason is the coercers needn't actually base their coercion method on established fact and truth, they can say to the voters, we've got a secret camera in the voting booth which will see where you put your X, even if there isn't one, and every time there's an election they can make a different story, this year it's a camera, next year it's cryptographic, and the year after is a magic box, or a talking dog that hides underneath the table that can hear which way they're voting. If the voters aren't educated enough to understand there's no such thing as a talking dog, then they're always going to lose this race, and the people who are attacking can try a different attack every year, and only reveal it afterwards. The people who are defending have got to try and educate the voters to call the bluff of every coercionist, so I think it's fundamentally a race that can't be won by the defenders. So I would say that adding complexity to improve coercion resistance is probably not very worthwhile. But improving accuracy, definitely.

Reply: Well we're targeting both.

Bruce Christianson: What we've got to do is to demonstrate that we have put countermeasures to these imaginary attacks into a protocol.

James Malcolm: Imaginary countermeasures?

Bruce Christianson: Well, possibly. But the real threat is that the imaginary attack is credible. That needs a real countermeasure[2].

Reply: This idea of giving voters encrypted receipts, and allowing them to check that they've been entered into the process, seems intellectually very appealing, but it may ultimately flounder on precisely that sort of issue, because you can have these con tricks where someone claims that they can read all of them.

Mike Bond: I suppose it's hearts and minds, because people could say, well it uses complicated cryptography, my vote must be secure, even if the coercer says that they can see it.

Bruce Christianson: It's as safe as your money is in an ATM. Is it possible to get schemes of this kind to scale linearly with fraud? What you hope is that buying ten votes is going to be ten times as difficult, or ten times as complicated as buying one vote, whereas with a lot of schemes there's a magic key, or some sort of magic point of failure, and once you've got that, you can commit fraud on an industrial scale.

[2] cf. LNCS 4631, p2.

I'm particularly interested in this choice of which side a human uses, because that seems like something that would be very hard to do automatically, and you don't want somebody to be able to do that automatically.

Reply: No, you very specifically don't want to make that automatic.

Bruce Christianson: As you say, there's a 50% chance of getting caught every time you do that. That seems like a nice property.

Reply: Yes. I think a nice feature of the two-sided scheme I suggested is this kind of fundamental symmetry between the two sides, so there shouldn't be a psychological bias in which choice, that there was to some extent with David Chaum's original scheme, because it involved visual crypto, and two transparent sheets overlaid, and so on, so there was definitely an upper side and a lower side, and one of the concerns there was whether there would be a distinct bias in voter choice between those two components. Hopefully with this scheme there will be less of a bias.

James Heather: Surely there's a bias whichever side is face up.

Reply: I don't quite know how you solve that, how we get it to print off vertically.

Michael Roe: Physical designs for voting schemes, it's like when we were doing lottery tickets, or raffle tickets, it gets printed and then tumbles down a chute, but hasn't had a chance to turn over a random number of times.

Reply: Yes, maybe there are things like that you can do.

Bruce Christianson: And it's possible that something like this could apply to lotteries and pools as well.

Would your system extend to not first past the post systems, to things like single transferable vote, and so forth?

Reply: Yes it does, certainly STV. You need a full permutation of a candidate list, so you can do that. If you're doing subset choices, that's slightly tricky because of the kind of problem that Matt was saying, you will be able to see how you selected, and that's difficult to avoid unless you have separate onions for each candidate.

Bruce Christianson: Or you could have random permutations between the two columns, you can have a yes column and a no column, but it's randomly swapped.

Reply: I guess you could do something like that, yes. Certainly you can do single transferable votes, and there are additional tricks: for single transferable vote, it's not a single cross, it's a list, a vector of values, and you can actually send these through the mix separately. One of the coercion problems with STV, particularly if you've got a large number of candidates, is that the low order values can be used and identified by a coercer, but if you send these through the mixer separately, you split them up, and so you avoid those kind of things. But yes, it does generalise.